Data ONTAP® 7.1
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 210-01077_B0
December 2005

Copyright and trademark information

Copyright information

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Portions Copyright © 2001, Sitraka Inc.

Portions Copyright © 2001, iAnywhere Solutions

Portions Copyright © 2001, i-net software GmbH

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Trademark information

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

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

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

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

Network Appliance NetCache is certified RealSystem compatible.

Table of Contents

Preface xi

Chapter 1 **Introducing NetApp Storage Systems** 1

About storage systems 2

About storage system components 4

About Data ONTAP 8

Chapter 2 **Interfacing with Data ONTAP** 15

How you administer a storage system 16

Using the command line interface 19

Using Data ONTAP commands at different privilege levels 22

Chapter 3 **Accessing the Storage System** 25

Access methods 26

Accessing a storage system from the console 29

Accessing a storage system using a Telnet session 33

Accessing a storage system using a Remote Shell connection 37

Accessing a storage system using the HTTP/Web interface, FilerView 40

Managing access from administration hosts 45

Controlling storage system access 49

Chapter 4 **Understanding the Root Volume** 55

Root volume size and space guarantee requirements 56

Default directories in the root volume 58

Contents of the /etc directory 59

Accessing the default directories on the storage system 61

Editing configuration files 63

Changing the root volume 65

Chapter 5 **Starting and Stopping the Storage System** 67

Booting the storage system 68
Booting from the storage system boot diskette set 70
Booting the F87 storage system 73
Booting the FAS200 and FAS3000 series storage systems 75
Booting the F800 and FAS900 series storage systems 85

Using storage systems as netboot servers 88

Rebooting the storage system 91

Halting the storage system 93

Chapter 6 **Managing Administrator Access** 95

Using administrator accounts to control administrative access 96
About managing administrator access 97
Managing users 100
Managing groups 105
Managing roles 108
Listing and deleting users, groups and roles 114
Administrative user creation examples 120

Managing passwords 122
Managing passwords for security 123
Changing passwords 125
Managing password rules 127

Chapter 7 **Performing General System Maintenance** 131

Aggregate Snapshot copy management 132

Managing licenses 136

Setting the system date and time 138

Synchronizing the system time 140

Understanding core files 143

Configuring message logging 145

Configuring audit logging 149

Configuring storage system startup 151
About the /etc/rc file 152
Editing the storage system's boot configuration file 154
Recovering from /etc/rc errors 155

Storage system configuration backup and cloning 156

UPS management 159

Chapter 8 **Using AutoSupport** 163

Learning about AutoSupport 164

Configuring AutoSupport. 166

Troubleshooting AutoSupport 172

Interpreting AutoSupport messages 175
About AutoSupport events 176
Contents of AutoSupport messages 177

Chapter 9 **Using SecureAdmin** 181

Learning about SecureAdmin. 182
How SecureAdmin uses SSH 183
How SecureAdmin uses SSL. 186

Managing SSH for SecureAdmin. 187
Setting up and starting SSH 188
Reinitializing SSH 190
Stopping or starting SSH service. 192
Setting up public key-based authentication 193

Managing SSL for SecureAdmin 197

General SecureAdmin administration 202

Chapter 10 **Managing Remotely with the RLM** 203

About the Remote LAN Module (RLM) 204

Configuring the RLM. 207

Managing the RLM with Data ONTAP 213

Logging in to the RLM 215

Managing the storage system with the RLM 218

Displaying storage system and RLM information 224

Comparing Data ONTAP and RLM commands 230

Troubleshooting the storage system with the RLM 233

Updating the RLM firmware 234

Troubleshooting RLM problems 238

Chapter 11 **System Information and Performance** 245

Storage system configuration information 246

Storage information. 249

Aggregate information 251

Volume information 253

File statistics for volumes. 255

About the filestats command 256

Options for the filestats command 259

Environmental status information 262

Fibre Channel information 265

Getting storage system information using the stats command 266

Getting system information using perfmon. 277

Improving storage system performance 278

Chapter 12 **Troubleshooting Tools** 283

Storage system panics 284

Error messages 285

Other troubleshooting resources 287

Glossary 289

Index 299

Preface

Introduction

This guide describes how to configure, operate, and manage storage systems that run Data ONTAP® 7.1 software. It covers F800 series storage systems, FAS storage systems and NearStore® storage 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 storage system's clients. It also assumes that you are familiar with how to configure the storage system 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® storage products (filers, FAS storage systems, 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 storage system through a session using Telnet, rsh, HTTP, Secure Socket Shell (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 storage system efficiently.

Attention

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

Introducing NetApp Storage 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 7.1 library, such as an aggregate, a FlexVol™ volume, a FlexClone™ volume (a FlexVol 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 storage systems" on page 2
- "About storage system components" on page 4
- "About Data ONTAP" on page 8

About storage systems

What a storage system is

A 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. Storage systems provide a modular hardware architecture running the Data ONTAP™ operating system and WAFL® (Write Anywhere File Layout) 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 V-Series systems provide storage for data. NetApp V-Series systems fulfill client requests from Logical Unit Numbers (LUNs) on the backend storage subsystems. For more information about V-Series systems, see the V-Series *Software Setup, Installation, and Management Guide*.

Data ONTAP is the operating system for all NetApp storage 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), and for storage systems with a Remote LAN Module (RLM) installed, through the RLM Ethernet connection to the system console. For more information about Data ONTAP, see “About Data ONTAP” on page 8.

Components of a storage system

A storage system consists of the following components.

Component	Function
Storage system main unit, or chassis, also known as the storage engine	The hardware device that receives and sends data. This unit also houses the storage system components and detects and gathers information about the hardware and the hardware configuration, the storage system components, operational status, hardware failures, and error conditions. For information about how to view this information, see "Environmental status information" on page 262. For information about environmental error codes, see the *Diagnostics Guide* on the NetApp on the Web™ (NOW) site at http://now.netapp.com.
Disk shelves	Containers, or *device carriers*, that hold disks and associated hardware (such as power supplies, connectivity, and cabling) that are connected to the main unit of the storage systems. For more information, see "Disk shelves and disks" on page 7. **Note** For V-Series systems, see the documentation for your storage subsystem for information about disks.

About storage system components

Types of components

The storage system has the following components:

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

Internal components

The following internal components enable the storage system to function.

Component	Description
System board	The system board is also referred to as the main board of the storage 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 storage 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 storage system in an emergency.

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

Component	Description
RLM (Remote LAN Module) (not available with all storage systems)	The RLM provides remote platform management capabilities for the storage system. It allows you to remotely access the storage system console over a network, and turn the storage system power on or off regardless of the operating state of the storage system. The RLM monitors and maintains hardware event logs for the storage system, and it generate alerts based on system status. For more information, see "About the Remote LAN Module (RLM)" on page 204.

Slots and ports

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

Component	Description
Slots	The storage 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 storage 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 V-Series systems, see the appropriate V-Series integration guide, the V-Series *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 storage system.

Storage 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. Storage systems can be fabric-attached, network-attached, or direct-attached to support NFS, CIFS, HTTP, and FTP (File Transfer Protocol) for file access, and Internet SCSI (iSCSI) for block-storage access, all over TCP/IP, as well as SCSI over Fibre Channel Protocol (FCP) for block-storage access, depending on your specific data storage and data management needs.

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

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

- 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
- 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*.

Data storage management: Data ONTAP stores data on disks in disk shelves connected to storage systems. 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*.

For SharedStorage™ storage systems, you can attach two to four systems to a common set of disk drives. For information, see the *Storage Management Guide*.

Note

For information about managing disks connected to V-Series systems, see the V-Series *Software Setup, Installation, and Management 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 your protocol.

When Data ONTAP is installed on an storage system 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 55.

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 Snapshot™ copies and makes them available so that users can access deleted or overwritten files. Snapshot copies are read-only copies of the entire file system.

  Snapshot copy is 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:

- Snapshot copy
- Asynchronous mirroring
- Synchronous mirroring
- Backup to tape
- Aggregate copy
- Volume copy
- FlexClone volume

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

Feature	Description
`aggr copy`	This is fast block copy of data stored in aggregates; it enables you to copy blocks of stored system data from one aggregate to another. For information about aggregates and `aggr copy`, see the *Storage Management Guide*.

Feature	Description
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 *Cluster Installation and Management Guide.*
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.*
SnapLock™ software (license required)	SnapLock provides an alternative to traditional optical WORM (write-once-read-many) storage systems for nonrewritable data. For information about SnapLock, see the *Data Protection Online Backup and Recovery Guide.*
SnapMirror® software (license required)	System-to-system Snapshot copy mirroring enables you to mirror Snapshot copies on one storage system to a partner system. Should the original storage system be disabled, this ensures quick restoration of data from the point of the last Snapshot copy. For information about SnapMirror, 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 Snapshot copies on an entire volume. For information about SnapRestore, see the *Data Protection Online Backup and Recovery Guide*.
Snapshot copy software	Snapshot copy software allows you to create manual or automatically scheduled multiple backups (or Snapshot copies) 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 Snapshot copy, see the *Data Protection Online Backup and Recovery Guide*.
SnapVault® software (license required)	SnapVault combines Snapshot copy schedules and Qtree SnapMirror (QSM) to provide disk-based data protection for NetApp storage systems. You can also install the Open Systems SnapVault agent on non-NetApp storage systems. This allows SnapVault to back up and restore data to those systems also. Using SnapVault, you can periodically replicate selected Snapshot copies from multiple client NetApp storage systems to a common Snapshot copy on the SnapVault server. The Snapshot copies 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
SyncMirror® (cluster configuration required)	The SyncMirror software performs real-time RAID-level—that is, RAID4 or RAID-DP™ (RAID double-parity)—mirroring of data to two separate plexes that are physically connected to the same storage system head. If there is an unrecoverable disk error on one plex, the storage 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*.
Tape backup and restore	Tape backup `dump` and `restore` commands enable you to back up system or SnapVault Snapshot copies to tape. Because the Snapshot copy rather than the active file system is backed up to tape, the storage 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*.
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.*
`vol copy`	This is fast block copy of data stored in volumes; it enables you to 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*.

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

- Manage network connections
- Manage adapters

- Manage protocols
- Configure pairs of storage systems into clusters for failover
- Configure SharedStorage storage systems into a community
- Manage storage
- Dump data to tape and restore it to the storage 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 storage system problems to technical support and up to five designated recipients. For more information about AutoSupport, see “Using AutoSupport” on page 163.

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 storage system" on page 16
- "Using the command line interface" on page 19
- "Using Data ONTAP commands at different privilege levels" on page 22

How you administer a storage system

Ways to administer a storage system

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

- Command execution through the storage system's command line interface (CLI)
- Command execution through the RLM redirection feature (not available on all platforms)
- 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 storage system's CLI

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

You can access the storage system's command line from

- A serial terminal connected to the console port of the storage system
- An Ethernet connection to an RLM installed in the storage system (not available for all platforms)
- A Telnet session to the storage 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 the RLM

You can use the redirection feature of the RLM to remotely execute all Data ONTAP administrative commands. For more information about this feature, see Chapter 10, "Managing Remotely with the RLM," on page 203.

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 storage 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 storage 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 58.

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 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 storage systems. The kit contains libraries, code samples and bindings in Java, C, and Perl for the new ONTAPI™ programming interface set. A NetApp storage system simulator which runs on Linux or Solaris, that simulates the NetApp storage 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 storage system command line:

```
toaster> 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 22.

Command line man pages

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

For more information, see the 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 storage 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 technical support.

Attention

Using advanced commands can be dangerous. Using some advanced commands without consulting 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 storage 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 [-q] [admin | advanced]
	admin sets the privilege level to administrative.
	advanced sets the privilege level to advanced.
	-q enables quiet mode. It suppresses the warning that normally appears when you set the privilege level to advanced.
	Note
	If no argument is given, the default, admin, is applied.
	Example: Assuming the name of the storage system is *sys1*, the storage system prompt is sys1>, as shown in the following example.
	`sys1> priv set advanced`
	Result: The following message is displayed, followed by the advanced mode storage system prompt.
	`Warning: These advanced commands are potentially dangerous; use them only when directed to do so by Network Appliance personnel.`
	`sys1*>`

Accessing the Storage System 3

About this chapter

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

Topics in this chapter

This chapter discusses the following topics:

- "Access methods" on page 26
- "Accessing a storage system from the console" on page 29
- "Accessing a storage system using a Telnet session" on page 33
- "Accessing a storage system using a Remote Shell connection" on page 37
- "Accessing a storage system using the HTTP/Web interface, FilerView" on page 40
- "Managing access from administration hosts" on page 45
- "Controlling storage system access" on page 49

Access methods

What type of access is needed

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

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

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

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

- From a console that is attached by a cable to the storage system's serial port
- From the Ethernet network interface card (NIC) that is preinstalled in the storage system. Use this card to connect to a TCP/IP network to administer the storage 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 (required for connecting the host to storage systems with an RLM installed)

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

You can directly access a storage 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 storage system's serial port
- From the Ethernet network interface card (NIC) that is preinstalled in the storage system. Use this card to connect to a TCP/IP network to administer the storage system, as well as to store and retrieve data:
  - 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 CIFS client by using a secure shell client application, such as SSH, OpenSSH for UNIX hosts or PuTTY for Windows hosts (required for connecting the host to storage systems with an RLM installed and using the console redirection feature)
- From an NFS client or CIFS client by using a Web browser and the DataFabric Manager interface (a DataFabric Manager 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 a console session

You use the `telnet.distinct.enable` option to control whether the console session is shared with a Telnet or an SSH-interactive session at the same time or the console session is a distinct user environment separate from Telnet and SSH-interactive sessions. However, if you have the RLM installed in your storage system, the console session is always shared with the RLM session, regardless of the `telnet.distinct.enable` option setting.

The default setting for the `telnet.distinct.enable` option is Off. This causes the console session to share a Telnet or SSH-interactive session. The `telnet.distinct.enable` option also determines when the autologout program logs you out of a Telnet session.

Note

An SSH-interactive session is initiated by opening the session without entering a command. For example, you would enter the following command: `ssh toaster -l root:""` (instead of `ssh filer -l root:""` *`command`*, which initiates a non-interactive session).

Rules that apply to console, Telnet, and SSH-interactive sessions

The following rules apply to console, Telnet, and SSH-interactive sessions regardless of the `telnet.distinct.enable` option setting:

- You cannot open more than one Telnet or SSH-interactive session at a time.
- You cannot open a Telnet and an SSH-interactive session at the same time.

Note

You can prevent commands from being aborted at the console or through a Telnet or SSH session by using the `rsh` command to initiate commands from an administration host.

Sharing the console session: If the `telnet.distinct.enable` option is set to Off, the console shares a session with a Telnet or SSH-interactive session, and the following rules apply:

- Commands typed at either the console or the Telnet or SSH-interactive session are echoed to the other location.
- Pressing Ctrl-C aborts the current command regardless of where the command was entered.
- Messages are displayed at both locations.
- Auditlog entries identify all console commands as "console shell," as shown in the following example:

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

- Auditlog entries identify all Telnet and SSH-interactive commands as "telnet shell."
- The autologout program logs the user out of the Telnet session after the number of minutes specified by the `autologout.telnet.timeout` option has elapsed. The timeout counter starts after the Enter or Return key is pressed. For example, if the `autologout.telnet.timeout` option is set to ten minutes, every time you press the Enter key, the timeout counter starts counting. If ten minutes elapse before you press the Enter key again, the autologout program logs you out.

Not sharing the console session: If the `telnet.distinct.enable` option is On, the console session has a distinct user environment and the following rules apply:

- 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 all console, Telnet, and SSH-interactive commands as "console shell."
- The autologout program logs the user out of the Telnet session after the number of minutes specified by the `autologout.telnet.timeout` option has elapsed. The timeout counter starts after the command is executed.

Accessing a storage system from the console

Using the serial port

To access a storage 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 storage system responds with the login or password prompt.
2	If the storage system displays the login prompt, do one of the following: ◆ To access the storage system with the system account, enter the following account name: `root` ◆ To access the storage system with an alternative administrative user account, enter the following: ***`username`*** *username* is the administrative user account. **Result:** The storage 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.

Step	Action

4	When you see the 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.

Using the RLM to remotely access the system console

The RLM is not available on all storage systems. For information about the RLM, see Chapter 10, "Managing Remotely with the RLM," on page 203.

You must use an account that has an assigned role as admin or root to access a system with the RLM.

To remotely access a system console by using the RLM system console redirection feature, complete the following steps.

Step	Action
1	From a UNIX or Windows administration host, log in to the RLM console. For information about the procedure, see "Logging in to the RLM" on page 215. **Result:** The storage system responds with the RLM prompt.

Step	Action
2	Enter the following command: `RLM toaster>` **`system console`** If the storage system displays the login prompt, do one of the following: ◆ To access the storage system with the system root account, enter the following account name: **`naroot`** ◆ To access the storage system with an alternative administrative user account, enter the following account name: ***`username`*** *username* is the administrative user account. **Result:** If the account requires a password, you are prompted for it.
3	Enter the password for the root or administrative user account. If no password is defined for the account, press Enter.

4	When you see the storage system prompt followed by a system message, press Enter to get to the storage system prompt.
	Example:
	`toaster> Tue Mar 1 15:19:39 PDI [toaster: ssh_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.

5	To exit the console redirection session and return to the RLM prompt, press Ctrl-D.

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 storage system using a Telnet session

Requirements for using a Telnet session

The following requirements must be met before you can connect to a storage 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 storage system allows Telnet access. For more information, see the na_options(1) and na_protocolaccess(1) 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 the console session" on page 28.

Using a Telnet session

To access a storage 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 storage system using its name.
3	If the storage system displays the login prompt, do one of the following: ◆ To access the storage system with the system account, enter the following account name: **`root`** ◆ To access the storage system with an alternative administrative user account, enter the following: ***`username`*** *username* is the administrative user account. **Result:** The storage 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 storage system prompt followed by a system message, press Return to get to the storage 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

You can terminate a Telnet session from the Telnet session, from a console session, or from a Remote Shell connection.

To terminate a Telnet session with a storage system, complete the following step.

Step	Action
1	Press Ctrl-] or enter the following command to log out of the storage system at the system prompt or at the console: **`logout telnet`** or press Ctrl-D to close the Telnet session. **Note** If you are at a Remote Shell connection, enter the following command: **`rsh -l username:password hostname logout telnet`**

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 storage 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 the timeout period for Telnet sessions. 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 storage system using a Remote Shell connection

Remote Shell with trusted remote hosts supported

The storage 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 storage 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 storage 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 storage 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 storage system.

Note

To issue commands from a Remote Shell on a PC, you must always supply a user name for the PC in the storage 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 storage system, and you can have up to 4 concurrent rsh sessions running on each vFiler unit.

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
```

Attention

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 storage 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 storage system's /etc/hosts.equiv file and issue the command.

If you want to take advantage of role-based user authentication, use SSH. For more information, see "Setting up public key-based authentication" on page 193.

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 storage system from a UNIX client by using an rsh command

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

To access a storage 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 backslash (`\`). 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 storage 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 storage 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 storage 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 and Java support

FilerView requires a Netscape Navigator® 4.51 or later or Microsoft® Internet Explorer™ 4.0 or later browser. The browser must have Java and JavaScript enabled.

If your system has a newer version of Microsoft Windows, and it does not include Java support, you must download a Java run-time environment separately to ensure FilerView functions properly.

Note

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

Options that control access to FilerView

The following options control access to FilerView. For information about how to use these options, see the na_options man pages:

- `httpd.admin.access`
- `httpd.admin.enable`
- `httpd.admin.ssl.enable`

Accessing a storage system from a client by using FilerView

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

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

Step	Action
3	Click FilerView. **Result:** ◆ If the storage 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 Secure Sockets Layer (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 181 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 unit 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 storage system remotely than to go to the storage system's physical location.

What an administration host is

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

Administration host privileges

The storage 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 storage system root directory and edit configuration files from the administration host. ◆ Enter Data ONTAP commands by using a Remote Shell connection.

If the administration host is...	You can...
A CIFS client	Edit configuration files from any CIFS client as long as you connect to the storage 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 storage system.

Requirements for using an NFS client

If you plan to use an NFS client to manage the storage 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 storage 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"]

Examples:

```
client1 peter
client1 "joe smith"
```

Note

If you access the storage system using `rsh` from an administration host listed in the /etc/hosts.equiv file, you have root privileges because this access method bypasses user authentication mechanisms. In addition, the /etc/auditlog program displays the user running the commands as root.

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 storage 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 storage system access

About this section

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

This section contains the following topics:

- "Controlling Telnet access using host names" on page 49
- "Controlling Telnet access using host names, IP addresses, and network interface names" on page 50
- "Controlling Remote Shell access" on page 50
- "Controlling mount privilege" on page 51
- "Controlling file ownership change privileges" on page 52
- "Controlling anonymous CIFS share lookups" on page 52
- "Options that help maintain security" on page 53

Controlling Telnet access using host names

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

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

Step	Action
1	Access the storage 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 storage 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 storage 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 storage 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 storage 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 storage system, complete the following steps.

Step	Action
1	Access the storage 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 storage system security. For information about valid values for the `trusted.hosts`, `telnet.access`, and `rsh.access` options, see the na_protocolaccess man page.

Option	Description
`trusted.hosts`	Specifies up to five hosts that are allowed Telnet, RSH and administrative HTTP (FilerView) access to the storage system for administrative purposes. The default is set to an asterisk (`*`), which allows access to all storage systems. This value is ignored for Telnet access if the `telnet.access` option is set. It is also ignored for administrative HTTP access if the `httpd.admin.access` option is set.
`telnet.access`	Controls which hosts can access the storage system through a Telnet session for administrative purposes. You can restrict Remote Shell access to the storage system by specifying host names, IP addresses, or network interface names. If this value is set, the `trusted.hosts` option is ignored for Telnet.
`rsh.access`	Controls which hosts can access the storage system through a Remote Shell session for administrative purposes. You can restrict Remote Shell access to the storage system by specifying host names, IP addresses, or network interface names.
`nfs.mount_rootonly`	Controls whether the storage 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.

Option	Description
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 storage 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 storage system's root volume, which contains special directories and configuration files that help you administer your storage system.

About the root volume

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

For storage systems that have Data ONTAP 7.0 or later installed at the factory, the root volume is a FlexVol volume. 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:

- "Root volume size and space guarantee requirements" on page 56
- "Default directories in the root volume" on page 58
- "Contents of the /etc directory" on page 59
- "Accessing the default directories on the storage system" on page 61
- "Editing configuration files" on page 63
- "Changing the root volume" on page 65

Root volume size and space guarantee requirements

About the size requirement for root FlexVol volumes

In contrast with traditional volumes, it is possible to create a FlexVol volume that is too small to be used as the root volume. The root volume must have enough space to contain system files, log files, and core files. If a system problem occurs, these files are needed to provide technical support.

Data ONTAP prevents you from setting the root option on a FlexVol volume that is smaller than the minimum root volume size for your storage system model, and prevents you from resizing the root volume below the minimum allowed size.

Minimum size for root FlexVol volumes

The minimum size for a root FlexVol volume depends on your storage system model. The following table lists the minimum allowed size for root volumes.

Storage system model	Minimum root FlexVol volume size
FAS250	9 GB
FAS270, FAS270c	10 GB
FAS920, FAS920c	12 GB
FAS940, FAS940c	14 GB
FAS960, FAS960c	19 GB
FAS980, FAS980c	23 GB
FAS3020, FAS3020c	12 GB
FAS3050, FAS3050c	16 GB
R100-12TB	13 GB
R100-24TB	19 GB
R100-48TB	30 GB
R100-96TB	53 GB
R150	19 GB
R200	19 GB

Root FlexVol volumes and space guarantees

A root FlexVol volume must have a space guarantee of `volume`. Data ONTAP prevents you from setting the root option on a FlexVol volume that does not have a space guarantee of `volume`, and prevents you from changing the space guarantee for the root volume.

Default directories in the root volume

What the default directories are

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

- The /etc directory, which contains configuration files that the storage 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 storage 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 storage system configuration files, executables required to boot the system, and some log files.

Attention

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

About the configuration files

Some of the configuration files in the /etc directory can be edited to affect the behavior of the storage system. If a configuration file can be edited by the system administrator, it is described in Section 5 of the man pages for your storage system. To edit a configuration file, use an editor on your administration host.

For more information about the /etc/rc file, see "Configuring storage system startup" on page 151. For more information about the quotas file, see the chapter on quotas in the *Storage Administration Guide*. For more information about other editable configuration files, see the man pages.

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 storage 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 storage 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 storage system. **Result:** You have access to the storage 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 storage system. **Result:** You have access to the storage 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 storage 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 storage system. **Note** You can also browse the Network Neighborhood to locate the storage 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 storage 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 storage system. You must use an editor from an NFS client or a CIFS client to edit storage 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 storage 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 storage system console, a Telnet client, or the `rsh` command to issue the following command to the storage system: **`exportfs`**
4	Mount the storage system root volume on the client.
5	From the client, use a text editor to edit the files in the /etc directory.

Editing configuration files from a CIFS client

After `setup` finishes, the default /etc/passwd and /etc/group files on the root volume are set up to enable you to share files on the storage system as Administrator. The storage system root directory is shared automatically as C$. The Administrator account 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 storage system as Administrator.
2	Display the contents of the storage 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 storage system's root volume, perhaps to migrate your root volume from a traditional volume to a FlexVol volume. Every storage system must have a root volume; therefore, you must always have one volume designated as the root volume.

Changing the root volume of a storage system

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

Note

To change your root volume from a traditional volume to a FlexVol volume or from a FlexVol volume 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 storage 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 Management Guide*. **Note** The required minimum size for the root volume varies depending on the storage system model. If the volume is too small to become the new root volume, you are prevented from setting the root option. For more information, see "Root volume size and space guarantee requirements" on page 56.
2	Using `ndmpcopy`, copy the /etc directory and all of its subdirectories from the current root volume to the new root volume. 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 storage system: **`reboot`** **Result:** When the storage system finishes rebooting, the root volume is changed to the specified volume.
5	Update the `httpd.rootdir` option to point to the new root 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 storage system reboot, there is also an *aggregate* `root` option. The aggregate `root` option is used only when, for some reason, the storage system cannot determine which volume to use as the root volume.

Note

You are advised not to change the value of the `root` option for any aggregate unless instructed to do so by technical support.

Starting and Stopping the Storage System 5

About this chapter

This chapter describes how to boot a storage system 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 storage system" on page 68
- "Using storage systems as netboot servers" on page 88
- "Rebooting the storage system" on page 91
- "Halting the storage system" on page 93

Booting the storage system

Where the storage system boots from

Storage 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 storage system from a version of Data ONTAP earlier than 7.0, your storage 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 storage system, see the *Upgrade Guide*.

Booting the storage system from the system or firmware prompt

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

- The system prompt (`toaster>`); if so, see "Rebooting the storage system" on page 91."
- The firmware prompt (`ok`)

To boot the storage system, complete the following step.

Step	Action
1	Enter the following command: **`boot`** **Note** If you are booting the storage 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 storage systems. For information about how your storage system boots, see the hardware and service guide for your storage system.

- "Booting from the storage system boot diskette set" on page 70
- "Booting the F87 storage system" on page 73
- "Booting the FAS200 and FAS3000 series storage systems" on page 75
- "Booting the F800 and FAS900 series storage systems" on page 85

Booting from the storage system boot diskette set

Storage systems that support floppy diskettes

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

Reasons to boot from diskette

You might need to boot the storage 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 storage 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.

Note

If you have upgraded your storage system, you must create new boot diskettes by downloading the appropriate Data ONTAP image from the NOW site. For information on how to perform this task, see the *Upgrade Guide*.

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 storage system is turned on</td><td>From the storage system console, enter the following command:
<code>reboot</code>
Result: The storage system begins the boot process.</td></tr>
<tr><td>Your storage system is turned off</td><td>Turn on your storage system.
Result: The storage system begins the boot process.</td></tr>
<tr><td>3</td><td colspan="2">When the storage system’s LCD prompts you, remove the diskette and insert the next diskettes into the storage system’s diskette drive when you are prompted.</td></tr>
<tr><td>4</td><td colspan="2">Press Enter.
Result: The storage system boots and, if it can, 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>
<tr><td>5</td><td colspan="2">Choose one of the boot types shown in the following table by entering the corresponding number.</td></tr>
</table>

Step	Action	
	Task	**Boot type**
	Run the storage system normally, but from a diskette.	Normal boot (1)
	Troubleshoot and repair configuration problems. **Note** Booting without /etc/rc causes the storage system to use only default options settings; disregard all options settings you put in /etc/rc; and disable some services, such as `syslog`.	Boot without /etc/rc (2)
	Change the password of the storage system.	Change password (3)
	Initialize all the disks and create a root traditional volume.	Initialize all disks (4)
	Initialize all the disks and create a root FlexVol volume.	Same as option 4, but create a root FlexVol volume. (4a)
	Perform aggregate and disk operations and display detailed information. Maintenance mode is special for the following reasons: ◆ Most normal functions, including file storage system operations, are disabled. ◆ A limited set of commands is available for diagnosing and repairing disk and aggregate and volume problems. ◆ Exit Maintenance mode with the `halt` command.	Maintenance mode boot (5)

Booting the F87 storage system

Boot options for the F87 storage system

The F87 storage 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 storage 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 storage 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 storage system from one of its SCSI disks

Note

You boot the storage system from a SCSI disk only when you install a new NVRAM adapter in the storage system. Do not use this option for normal operation of the storage 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 storage system.
	
	Result : The storage system console displays power-on self-test (POST) messages. After POST successfully executes, the storage 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 storage system boots the current release of Data ONTAP.

Booting the FAS200 and FAS3000 series storage systems

Boot options

The FAS200 series (FAS250, FAS270, and FAS270c) and FAS3000 storage 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 and FAS3000 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 storage 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 storage system. It provides a “known good” release from which you can boot the storage 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 storage system (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, enable netboot immediately for your storage system by placing a boot image on a local server. For more information, see "Preparing a storage system for netboot startup" on page 82.

- `boot_diags`

  Boots a Data ONTAP diagnostic kernel. For more information, see the *Diagnostics Guide*.

Note

Other boot options should be used only under the direction of technical staff.

Booting Data ONTAP at the firmware prompt

For FAS200 and FAS3000 series storage systems, 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 storage system console, enter the following command: **`halt`** **Result:** The storage 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.

Booting FAS3000 storage systems

The FAS3000 series storage systems are configured to boot from 256-MB PC CompactFlash cards. If you have an RLM installed, you can boot or reboot the storage system remotely.

If you reboot your storage system, it 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 FAS3000 series storage system, complete the following steps.

Step	Action		
	If...	**Without RLM, then...**	**With RLM, then...**

1	Your storage system is turned on	From the storage system console, enter the following command:
		`reboot`
		Result: The storage system begins the boot process.
			Open two RLM sessions. From one session, enter the following commands:
			`system console`
			`reboot`
			Result: The storage system begins the boot process.

	Your storage system is turned off	Turn on your storage system. **Result:** The storage system begins the boot process.	From the RLM session that is still running, enter the following command: `system power on` **Result:** The storage system begins the boot process.
2	You want the storage system to boot automatically in normal mode	Allow the storage system to reboot uninterrupted until the following prompt appears: `root logged in from console` You are done.	

<table>
<tr><th>Step</th><th colspan="3">Action</th></tr>
<tr><td></td><td>If...</td><td>Without RLM, then...</td><td>With RLM, then...</td></tr>
<tr><td>3</td><td>You want to select from a menu of alternative boot modes</td><td colspan="2">When prompted to do so, press Ctrl-C to display the boot menu.
Result: The storage 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>
<tr><td>4</td><td colspan="3">Choose one of the boot types shown in the following table by entering the corresponding number.</td></tr>
<tr><td></td><td>Task</td><td colspan="2">Boot type</td></tr>
<tr><td></td><td>Boot the storage system normally.</td><td colspan="2">Normal boot (1)</td></tr>
<tr><td></td><td>Troubleshoot and repair configuration problems.
Note
Booting without /etc/rc causes the storage 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 colspan="2">Boot without /etc/rc (2)</td></tr>
<tr><td></td><td>Change the password of the storage system.</td><td colspan="2">Change password (3)</td></tr>
</table>

Step	Action		
	If...	**Without RLM, then...**	**With RLM, then...**
	Initialize all the disks and create a root traditional volume.	Initialize all disks (4)	
	Initialize all the disks and create a root FlexVol volume.	Same as option 4, but create a root FlexVol volume. (4a)	
	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 storage system, enter **`boot`** after the firmware `ok` prompt.	Maintenance mode boot (5)	

Booting Data ONTAP remotely

For FAS3000 series storage systems with the RLM installed, you can boot Data ONTAP remotely. For more information about the RLM, see Chapter 10, “Managing Remotely with the RLM,” on page 203.

To boot Data ONTAP, complete the following steps.

Step	Action
1	Log in to the RLM by entering the following command at any UNIX or Windows host connected to the storage system (assuming the appropriate secure shell protocol application is installed). For example: `UNIX_host> ssh username@RLM_IP_address` or Windows host: Open a Windows session with PuTTY with the ssh setting selected, and the username and the RLM IP address in the appropriate fields. **Result:** The RLM prompt appears. **Example:** `RLM toaster>`
2	At the RLM prompt, enter the following command: `RLM toaster> system console` **Result:** The storage system prompt appears. **Example:** `toaster>`
3	At the storage system prompt, enter the following command: `toaster> halt` **Result:** The storage system console displays the `CFE>` firmware prompt.
4	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.

Recovering from a corrupted CompactFlash image

To recover from a corrupted CompactFlash image for a FAS3000 series storage system with RLM installed, complete the following steps.

Step	Action
1	Log in to the RLM by entering the following command at the administration host. `admin_host> ssh username@RLM_IP_address` **Result:** The RLM prompt appears. **Example:** `RLM toaster>`

2	Enter the following command:
	`RLM toaster> system reboot backup`
	Result: The following prompt is displayed:
	`This will cause a dirty shutdown of your appliance. Continue? [y/n]`

3	Enter `y` to continue. **Result:** The storage system shuts down abruptly. If the NVRAM contains data, the RED internal LED (seen through the face plate of the system) blinks. When the system is rebooted, the NVRAM automatically and transparently replays the data transactions.

Checking available Data ONTAP versions

By default, an FAS200 series storage 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 storage 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 storage system console, enter...
The current booted Data ONTAP version	`version` **Result:** The console displays the version number of Data ONTAP that is currently running.

To determine...	At the storage system console, enter...
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 storage system for netboot startup

To use the netboot option to start your storage 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/.

  Note
  You can also store the boot image on another storage system.

- Networking configured in the FAS200 series or FAS3000 firmware environment. You can use a Dynamic Host Configuration Protocol (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 storage system through the netboot option

To boot Data ONTAP from a remote image for FAS200 or FAS3000 series systems, complete the following steps.

Step	Action	
1	At the storage 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 storage system. *netmask* is the network mask of the storage system. *gateway* is the gateway for the storage 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 F800 and FAS900 series storage systems

Booting F800 and FAS900 storage systems with CompactFlash cards

F800 series and FAS900 series storage systems are configured to boot from 32-MB or 256-MB PC CompactFlash cards. If you have upgraded your F800 and FAS900 storage 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 storage system, it 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 storage system or an FAS900 series storage system, complete the following steps.

Step	Action	
1	**If....**	**Then...**
	Your storage system is turned on	From the storage system console, enter the following command: `reboot` **Result:** The storage system begins the boot process.
	Your storage system is turned off	Turn on your storage system. **Result:** The storage 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 storage system to boot automatically in normal mode</td><td>Allow the storage 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 storage system 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>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 storage system normally.</td><td>Normal boot (1)</td></tr>
<tr><td>Troubleshoot and repair configuration problems.
Note
Booting without /etc/rc causes the storage 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 storage system.</td><td>Change password (3)</td></tr>
<tr><td>Initialize all the disks and create a root traditional volume.</td><td>Initialize all disks (4)</td></tr>
<tr><td>Initialize all the disks and create a root FlexVol volume.</td><td>Same as option 4, but create a root FlexVol 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 storage system, enter <code>boot</code> after the firmware <code>ok</code> prompt.</td><td>Maintenance mode boot (5)</td></tr>
</table>
</td></tr>
</table>

Using storage systems as netboot servers

Netboot server requirements

You can configure a storage 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 storage system.
- The rest of your netboot-using environment to access the storage 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 storage 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 custom-mips in that directory, the boot image will be available for netboot startup at the following URL:

http://your.filer.com/na_admin/boot/custom-mips

Configuring HTTP services

To configure a storage system as an HTTP netboot server, complete the following steps.

Step	Action
1	Place boot programs in the storage 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 storage 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 storage 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 storage system

About rebooting

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

Rebooting the storage system from the system console

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

Step	Action
1	Send an advance warning to CIFS users to alert them to save their files and close any applications. **Attention** Never interrupt CIFS service by halting the storage 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.
2	At the storage system prompt, enter the following command: `toaster> reboot [-t minutes]` `-t minutes` is the amount of time that elapses before the reboot occurs.

Rebooting the storage system remotely

You can reboot a FAS3000 series storage system remotely if you are able to log in to the RLM installed in the storage system. To reboot the storage system remotely from the RLM, complete the following steps.

Step	Action
1	From the administration host, log in to the RLM. For more information, see "Logging in to the RLM" on page 215.

Step	Action
2	From the RLM prompt, enter the following command. `RLM toaster>` **`system console`**
3	From the storage system console prompt, enter the following command: `toaster>` **`reboot`**

Halting the storage 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 storage 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 storage 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 storage system, the users' applications are terminated and changes made to open files since the last save are lost.

Attention

Never interrupt CIFS service by halting the storage 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 storage system is built into Windows NT and Windows XP.

Halting the storage system

To halt the storage system, complete the following step.

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

Managing Administrator Access 6

About this chapter

This chapter describes how to control administrator access to your storage 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 96
- "Managing passwords" on page 122

Using administrator accounts to control administrative access

About this section

This section contains the following topics:

- "About managing administrator access" on page 97
- "Managing users" on page 100
- "Managing groups" on page 105
- "Managing roles" on page 108
- "Listing and deleting users, groups and roles" on page 114
- "Administrative user creation examples" on page 120

About managing administrator access

Reasons for creating administrator accounts

You can use the default system administration account, or root, for managing a storage 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 storage systems.
- You can limit an administrator's access to specific storage 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 storage system.

  The auditlog file keeps a record of all administrator operations performed on the storage 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 storage system.
- If a storage system 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 storage system. Terms used in this section include the following:

user: An account that is authenticated on the storage system. Users can be placed into storage system groups to grant them capabilities on the storage 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 storage system groups, thereby being granted capabilities on the storage system. This only works if CIFS has been set up on the storage system.

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 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 storage 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 storage 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.

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. Or, use `useradmin user modify` to modify the attributes of an existing user. **`useradmin user {add|modify} user_name [-c comments] [-n full_name] [-p password] -g group1[,group2,group3,..] [-m password_min_age] [-M password_max_age]`** *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 98. *comments* specifies a maximum 128-character comment which can be viewed through the `useradmin list` command. Comments cannot contain a colon character (:). *full_name* specifies the full name for the user. *password* is the password required of the specified administrative user (used only for rsh access). If the `security.passwd.rules.enable` option is set to On, the password must conform to the rules specified by the `security.passwd.rules.*` options. *group* is a predefined or customized group with roles assigned through the `useradmin group` command. For more information about creating custom groups, see "Assigning roles to groups by creating or modifying a group" on page 106. For a description of the predefined groups, see "Predefined groups" on page 105. *password_min_age* specifies the minimum number of days users can have a password before they can change it. The default value is 0. *password_max_age* specifies the maximum number of days users can have a password before they are required to change it. The default value is 4,294,967,295.

Step	Action
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 -n "Molly Mulberry" -c "Filer
administrator in Corp IT" -g Administrators
```

Granting access to Windows domain users

You can specify nonlocal administrative users to have administrative access to the storage system after authentication by a Windows Domain Controller, rather than by the storage system itself.

Note

Nonlocal users, or users added with the `useradmin domainuser add` command, cannot log in to a storage system using the login protocols (telnet, console, rsh, ssh or http-admin). Nonlocal 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 storage system group the nonlocal user has been placed in.

To grant a Windows domain user administrative access to an storage system, 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 storage 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 105. `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 100. **Example:** The following command 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 ageing characteristics of a user's password
- 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 by creating or modifying a group

To create or modify a group, giving that group the capabilities associated with one or more predefined or customized 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 98. **Note** Do not create groups with the same name as any of the Windows special groups or any existing users. For a list of the Windows special groups, see "Windows special groups" on page 99. *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 108. **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`**

Step	Action
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 99.

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 storage system using Telnet. ◆ `login-console`—gives the specified role the ability to log in to the storage system using the console. ◆ `login-rsh`—gives the specified role the ability to log in to the storage system using `rsh`. ◆ `login-ssh`—gives the specified role the ability to log in to the storage system using SSH. ◆ `login-http-admin`—gives the specified role the ability to log in to the storage 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 104. ◆ `security-priv-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 98. *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 storage system. 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 Password min/max age in days: 0/4294967295 Status: enabled
useradmin user list -g *grp_name*	Lists information for all users assigned to a specified group. For example: **> useradmin user list -g Admistrators** Name: administrator Info: Built-in account for administering the filer Rid: 500 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 storage system. 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.`
	`Rid: 544`
	`Roles: admin`
	`Allowed Capabilities: login-*,cli-*,api-*,security-*`

Command	Description

`useradmin role list`	Lists all the roles configured for this storage system. 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 storage system. 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 104.
`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 storage system, use the `useradmin user delete` command instead.
`useradmin group delete` *group_name*	Deletes the specified group from the storage system. **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 storage 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 storage system that are in local groups but do not have console access or any administrative capabilities on the storage 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 storage 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

About this section

This section describes managing passwords on the storage system. The following topics are included:

- "Managing passwords for security" on page 123
- "Changing passwords" on page 125
- "Managing password rules" on page 127

Managing passwords for security

How you can manage passwords for security

Data ONTAP provides several methods you can use to ensure the password policies for your storage systems meet your company's security requirements. The following list outlines these methods:

- Password rules

  Password rules enable you to specify rules for valid passwords. You use the `security.passwd.rules` options to specify password rules.

  For more information, see "Managing passwords" on page 127 or the na_options(1) man page.

- Password history

  Password history enables you to require users to rotate through a specified number of passwords, rather than simply using the same password every time. You use the `security.passwd.rules.history` option to specify password history. The default value is 0, which does not enforce this rule.

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

- Password expiration (maximum age)

  Password expiration enables you to require that users change their passwords before they are a specified number of days old. You use the `useradmin user add` or `modify` commands to set this value for individual users. The default value is 4,294,967,295.

  Note

  Before using password expiration, make sure your storage system time is set correctly. If you use password expiration before the date is set correctly, accounts could expire before or after the desired expiration date.

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

- Password minimum age

  Password minimum age prevents users from changing their passwords too quickly, thus cycling through their previous passwords too quickly. You use the `useradmin user add` or `modify` commands to set this value for individual users. The default value is 0, which does not enforce a minimum password age.

Note

Before using password minimum ages, make sure your storage system time is set correctly. Changing the system time after password minimum ages have been set can lead to unexpected results.

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

- Password lockout

  Password lockout enables you to lock users out after a specified number of unsuccessful login attempts. This is to prevent an unauthorized user from attempting to guess a password. You use the `security.passwd.lockout.numtries` option to specify password lockout. The default value is 0, which does not enforce this rule.

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

- Password reset requirement

  The password reset requirement enables you to require that all new users (except for root) reset their passwords when they log in for the first time. Users must also reset their passwords the first time they log in after another user has changed their password.

  You set the `security.passwd.firstlogin.enable` option to On to enable this requirement. The default value is Off.

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

Changing passwords

Changing the storage system password

To change the storage 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 storage system	1. Enter the following command: **`passwd`** 2. Enter the storage system account name: **`root`** 3. Enter the existing storage system password (not required if you are 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 storage 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 storage 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 storage system	Enter the following command: **`rsh` *`filer_name`* `passwd` *`new_password username`***

Managing password rules

Managing password rules

Data ONTAP provides the following set of options to control password rules:

Password rule option	Description
`security.passwd.firstlogin.enable {on|off}`	Specifies whether new users, and users logging in for the first time after another user has changed their password, must change their password. The default value for this option is Off. **Note** If you enable this option, you must ensure that all groups have the `login-telnet` and `cli-passwd*` capabilities. Users in groups that do not have these capabilities cannot log in to the storage system.
`security.passwd.lockout.numtries` *num*	Specifies the number of allowable login attempts before a user's account is disabled. The default value for this option is 4,294,967,295.
`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 is 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).

Password rule option	Description
`security.passwd.rules.everyone {on|off}`	Specifies whether a check for password composition is performed for the “root” and “Administrator” users. If the `security.passwd.rules.enable` option is set to Off, this option does not apply. The default value for this option is Off.
`security.passwd.rules.history` *num*	Specifies the number of previous passwords that are checked against a new password to disallow repeats. The default value for this option is 0, which means that repeat passwords are allowed.
`security.passwd.rules.maximum` *max_num*	Specifies the maximum number of characters a password can have. **Note** This option can be set to a value greater than 16, but a maximum of 16 characters are used to match the password. Users with passwords longer than 14 characters will not be able to log in via the Windows interfaces, so if you are using Windows, do not set this option higher than 14. The default value for this option is 256.
`security.passwd.rules.minimum` *min_num*	Specifies the minimum number of characters a password must have. The default value for this option is 8.
`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.
`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.

Password rule option	Description
`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 storage system.

Topics in this chapter

This chapter discusses the following topics:

- "Aggregate Snapshot copy management" on page 132
- "Managing licenses" on page 136
- "Setting the system date and time" on page 138
- "Synchronizing the system time" on page 140
- "Understanding core files" on page 143
- "Configuring message logging" on page 145
- "Configuring audit logging" on page 149
- "Configuring storage system startup" on page 151
- "Storage system configuration backup and cloning" on page 156
- "UPS management" on page 159

Aggregate Snapshot copy management

What an aggregate Snapshot copy is

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

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

What you can do with aggregate Snapshot copies

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

You use aggregate Snapshot copies in the following situations:

- If you are using MetroCluster or RAID SyncMirror and you need to break your mirror, you take an aggregate Snapshot copy 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 storage system, and you want to be able to restore the entire system state if the change produces unexpected results, you take an aggregate Snapshot copy before making the change.
- If the aggregate file system becomes inconsistent, aggregate Snapshot copies can be used by technical support to restore the file system to a consistent state.

How you manage aggregate Snapshot copies

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

How you create aggregate Snapshot copies

Usually, you do not need to create aggregate Snapshot copies manually. A schedule is automatically set up to generate new aggregate Snapshot copies periodically. In most cases, you should not need to change the aggregate Snapshot copy schedule.

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

About the aggregate Snapshot reserve

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

The default size of the aggregate Snapshot reserve is five percent of the aggregate size. For example, if the size of your aggregate is 500 GB, then 25 GB is set aside for aggregate Snapshot copies.

Note

Unlike volume Snapshot copies, aggregate Snapshot copies cannot consume any space outside of their Snapshot reserve.

About automatic aggregate Snapshot copy deletion

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

When an aggregate Snapshot copy 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 five percent is sufficient. However, if you find that aggregate Snapshot copies are being created and deleted often enough to affect system performance, then you should increase the aggregate Snapshot reserve.

Note

If you have automatic aggregate Snapshot copy creation enabled, you should not decrease the size of the aggregate Snapshot reserve below the default of five percent. If you need to reclaim the space being used for the aggregate Snapshot reserve, disable automatic aggregate Snapshot copy creation using the procedure outlined in "About disabling automatic aggregate Snapshot copy creation" on page 135.

How you manage automatic aggregate Snapshot copy deletion

In most cases you should leave automatic aggregate Snapshot copy deletion enabled. If this option is turned off for a particular aggregate, then every volume in that aggregate requires up to two 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 copy deletion temporarily. For example, if you want to create an aggregate Snapshot copy before breaking your RAID SyncMirror, you would want to make sure that Snapshot copy is not automatically deleted. To disable automatic aggregate Snapshot copy deletion, you use the `aggregate option` command.

For example, to turn off automatic aggregate Snapshot copy deletion 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 automatic aggregate Snapshot copy deletion, then space guarantees will be disabled for one or more of your volumes. For this reason, you should plan to re-enable automatic aggregate Snapshot copy deletion as quickly as possible.

About disabling automatic aggregate Snapshot copy creation

You can turn off automatic aggregate Snapshot copy creation for a particular aggregate, using the same `nosnap` option that you would for volume Snapshot copy. However, you are advised to leave automatic aggregate Snapshot copy creation enabled, in case you need any low-level file system repair.

Attention

Do not disable automatic aggregate Snapshot copy creation if you have a MetroCluster configuration, or if you are using RAID SyncMirror. Aggregate Snapshot copies are required for those configurations for correct system operation.

Disabling automatic aggregate Snapshot copy creation

To disable automatic aggregate Snapshot copy creation, and reclaim the free space being used for the aggregate Snapshot reserve, complete the following steps.

Step	Action
1	Disable automatic aggregate Snapshot copy creation by entering the following command: **`aggr options aggr_name nosnap on`** *aggr_name* is the name of the aggregate for which you want to disable automatic Snapshot copy creation.
2	Delete all Snapshot copies in the aggregate by entering the following command: **`snap delete -A -a aggr_name`**
3	Set the aggregate Snapshot reserve to 0 percent by entering the following command: **`snap reserve -A aggr_name 0`**

Managing licenses

About license codes

A license code is a string of characters, such as ABCDEFG, that is unique to a particular service. You receive license codes for every protocol and option, or service, that you purchase. Not all purchased license codes are installed on a storage 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 services at any time.

License code replacement

If you misplace a license code, you can contact technical support to obtain a copy.

License management tasks

You can perform the following tasks to manage licenses:

- Add licenses
- Display all services, including which licenses have been installed
- Delete licenses

Adding a license

To add the license for a service, complete the following step.

Step	Action
1	Enter the following command: **`license add <code1> <code2>...`** *code* is the license code provided to you by your sales person or technical support.

Displaying current license codes

To display licensing information for all services, 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 service, 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 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 storage system can service requests correctly. To keep your storage system's date and time correct automatically, see "Synchronizing the system time" on page 140.

Step	Action
1	Access the storage 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 storage system's date earlier when SnapMirror is running, Snapshot copies can appear out of sequence. When this occurs, SnapMirror assumes that the older Snapshot copy 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 storage 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 storage system synchronized with a time server. Using this feature is advised, because problems can occur when the storage system clock is inaccurate.

About time servers

To automatically keep your storage system time synchronized, you need the name of at least one time server. For best results, 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 storage system time synchronized with the time server, complete the following steps.

Step	Action
1	If the current time for the storage 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 138.

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 140. For more information about the `timed` options, see the na_options(1) man page or "About the timed options" on page 141.
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.	◆ *n*s ◆ *n*m ◆ *n*h	`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 storage system to panic, the system creates a core file that technical support can use to troubleshoot the problem. The storage 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, performs the following tasks:

- 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 storage 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 storage 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 storage system sends e-mail automatically to technical support upon each system reboot, if the AutoSupport feature is enabled and configured correctly. 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 technical support when your system creates a core file.

Configuring message logging

What message logging is

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

The level of information that the storage 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 storage 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 storage system" on page 61.

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 storage system daemons, such as `rshd`
`kern`	Messages from the storage 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 146.
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, an SSH 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 storage system" on page 61.

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 storage 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 152
- "Editing the storage system's boot configuration file" on page 154
- "Recovering from /etc/rc errors" on page 155

About the /etc/rc file

Startup commands are stored in /etc/rc file

The /etc/rc file contains commands that the storage 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 storage system to

- Communicate on your network
- Use the NIS and DNS services
- Save the core dump that might exist if the storage system panicked 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 storage 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 storage 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 storage 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 storage 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 143.

Editing the storage system's boot configuration file

What the storage system's boot configuration file contains

The storage 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 storage system's boot configuration file

To edit the storage 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 storage 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 storage 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 storage system becomes inaccessible to the administration host

The storage 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 storage system is inaccessible because it is not on the network.
- You improperly export storage system directories to the NFS client that is the administration host. The storage 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.	
	If you are in...	**Then...**
	An NFS environment	Enter the `exportfs` command to export the storage system root directory to the administration host.
	A CIFS environment	Add a share to the storage system root directory.
2	Edit the storage system /etc/rc file from the administration host.	
3	Reboot the storage system.	
4	If the changes do not correct the problem, repeat Step 1 through Step 3.	

Storage system configuration backup and cloning

About storage system configuration backup and cloning

Storage 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 storage system configuration in case of disasters or emergencies. Configuration cloning enables you to clone the configuration of an existing storage system to a new system.

Backing up a storage system configuration

When you use the `config` command to back up a storage 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 storage 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 storage system-specific configuration to the file /etc/configs/08_02_2004.
	`config dump 08_02_2004`
	Example 2: The following is an example of the `config dump` command with a directory that you specify.
	`config dump /home/users/08_02_2004`

Cloning a storage system configuration

To clone the configuration of one storage 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 storage system from which you want to clone the configuration. *username* is the login name of an administrative user on the remote storage 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 storage system toaster. `toaster> config clone tpubs-dot root:hello`

About restoring a storage system configuration

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

Restoring a storage system configuration

To restore storage 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 storage 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_2004`

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

Comparing storage system configurations and backup configuration files

To compare a storage 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 storage system's current configuration with the configuration information in the backup file.
	`config diff 11_15_2004`
	Example 2: The following example compares the configuration information in two backup files.
	`config diff -o diff.txt 11_05_2004 11_15_2004`

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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 five minutes for `warningtime` and one minute for `criticaltime`. However, you are advised to make sure that 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 storage system will not be able to shut down cleanly.

- Storage 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

You are strongly advised to contact 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 storage system to 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 164
- "Configuring AutoSupport" on page 166
- "Troubleshooting AutoSupport" on page 172
- "Interpreting AutoSupport messages" on page 175

Learning about AutoSupport

About AutoSupport

The `autosupport` daemon monitors the storage system's operations and sends automatic messages to technical support to alert it to potential system problems. If necessary, 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 storage system.
- AutoSupport messages are generated
  - When events occur on the storage system that require corrective action from the system administrator or technical support
  - When the storage system reboots
  - When you initiate a test message using the `autosupport.doit` option
  - Once a week, early Sunday morning, at approximately midnight

    There are two AutoSupport messages generated at this time. One, the weekly AutoSupport message, provides the same system information as regular AutoSupport messages. The other, the performance AutoSupport message, provides technical support with comprehensive performance information about your storage system for the preceding week. The performance message can be quite large, so by default, it is sent only to technical support.
- The system can send AutoSupport messages by SMTP, HTTP, or HTTPS (the Secure Sockets Layer 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 storage 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 you are strongly advised to leave it enabled. Enabling AutoSupport can significantly speed problem determination and resolution should a problem occur on your storage system.

AutoSupport transport protocols

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

- HTTPS
- HTTP
- SMTP

The default transport protocol is HTTPS. Because SMTP can introduce limitations on message length and line length, you are strongly advised to 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 storage 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 storage system uses the mail host's mail server to send periodic e-mail messages automatically to 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 Multipurpose Internet Mail Extensions (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 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 NetApp, 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 172.
`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 storage 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. 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 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. The default is `https`.
`autosupport.support.url`	Indicates where AutoSupport notifications 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 166: **`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.

Note

If you use the keyword TEST in the message, you receive a return message indicating that the AutoSupport process is working correctly. For more details, see the information on the `autosupport.doit` command in "AutoSupport options" on page 166.

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 storage 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 storage system can talk to by entering the following command on the storage 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.yourco.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.

Step	Action
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 WARNING 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.

Step	Action
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 advised, as it may impact the quality of your technical support.

Interpreting AutoSupport messages

About AutoSupport messages

AutoSupport messages can help you understand the status and operations of your storage 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 176
- "Contents of AutoSupport messages" on page 177

About AutoSupport events

Event and log-level descriptions

The storage system sends AutoSupport messages to technical support about your system after any of several events. The AutoSupport message includes a log level that indicates the priority assignment from 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 find the Message Matrices page.
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 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 that command.

Each AutoSupport message contains the following types of information.

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

Note

For more information about setting AutoSupport options, see "AutoSupport options" on page 166. 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 storage system.

Topics in this chapter

This chapter discusses the following topics:

- "Learning about SecureAdmin" on page 182
- "Managing SSH for SecureAdmin" on page 187
- "Managing SSL for SecureAdmin" on page 197
- "General SecureAdmin administration" on page 202

Learning about SecureAdmin

How SecureAdmin improves security

SecureAdmin makes it very difficult for someone to intercept a storage 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 storage 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 183
- "How SecureAdmin uses SSL" on page 186

How SecureAdmin uses SSH

How SSH improves security

SSH improves security by providing a means for a storage system to authenticate the client and by generating a session key that encrypts data sent between the client and storage 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

Host 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.

Server key: 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 storage system reboots

Session key: SSH uses the session key to encrypt data sent between the client and storage 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 storage system, where it is decrypted using the host and server keys. After the session key is decrypted, the client and storage 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 storage system and client.

Stage	What the client does	What the storage system does
1	The client sends an SSH request to the storage system.	The storage system receives the SSH request from the client.
2		The storage 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 storage 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 storage system.	
6		The storage 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 storage 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 storage 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 storage system enables the SSL capabilities between system and browser.

Unlike using FilerView to send the storage 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 storage 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 storage system.

Ways to manage SSH

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

- "Setting up and starting SSH" on page 188
- "Reinitializing SSH" on page 190
- "Stopping or starting SSH service" on page 192
- "Setting up public key-based authentication" on page 193

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 storage 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. You can 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 storage 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 storage 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 storage 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 storage 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 storage system root volume.

The following commands append the public key to the /etc/sshd/*user_name*/.ssh/authorized_keys file on storage 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 195 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 storage 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 storage system root volume.

The following commands append the public keys to the /etc/sshd/*user_name*/.ssh/authorized_keys file on storage 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 storage 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	Action
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">Action</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 200 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 199 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 storage 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 storage 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 200 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 storage 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 storage system, because the domain name recorded in the certificate becomes obsolete. The result is that the connection is still encrypted, but the storage 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 storage 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 197.

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 storage 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`

Managing Remotely with the RLM 10

About this chapter

This chapter describes what the Remote LAN Module (RLM) is, what it does, and how to remotely manage a storage system with it.

Topics in this chapter

This chapter discusses the following topics:

- "About the Remote LAN Module (RLM)" on page 204
- "Configuring the RLM" on page 207
- "Managing the RLM with Data ONTAP" on page 213
- "Logging in to the RLM" on page 215
- "Managing the storage system with the RLM" on page 218
- "Displaying storage system and RLM information" on page 224
- "Comparing Data ONTAP and RLM commands" on page 230
- "Troubleshooting the storage system with the RLM" on page 233
- "Updating the RLM firmware" on page 234
- "Troubleshooting RLM problems" on page 238

About the Remote LAN Module (RLM)

What the RLM is

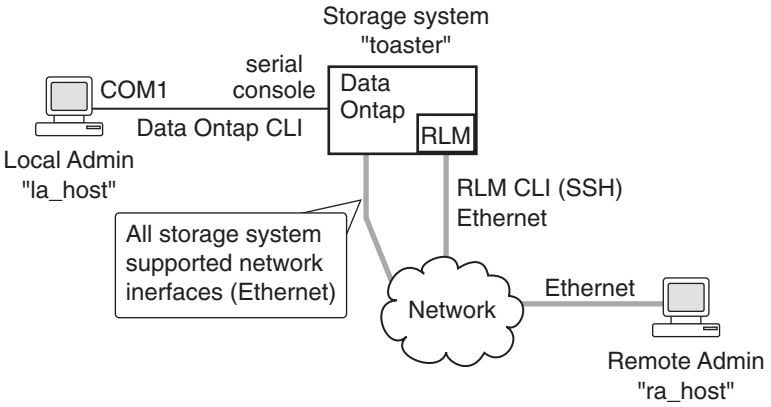
The RLM is a remote management card that can be installed inFAS3000 series storage systems to provide remote platform management capabilities, including remote access, monitoring, troubleshooting, logging and alerting features.

The RLM stays operational regardless of the operating state of the storage system. It is powered by a standby voltage, which is available as long as the storage system has input power to at least one of the storage system's power supplies.

The FAS3000 series storage systems provide an Ethernet interface for connecting to the RLM. If the RLM is not pre-installed in your storage system, see the *Installing or Replacing a Remote LAN Module* flyer for instructions on how to cable your storage system to the RLM. The flyer is shipped with the RLM and is also available on the NOW Web site at http://now.netapp.com/.

The following diagram illustrates how you can access the storage system and the RLM.

- Without the RLM, you can *locally* access the storage system through the serial console or from an Ethernet connection using any supported network interface. You use the Data ONTAP CLI to administer the storage system.
- With the RLM, you can *remotely* access the storage system through the serial console. The RLM is directly connected to the storage system through the serial console. You use the Data ONTAP CLI to administer the storage system and the RLM.
- With the RLM, you can also access the storage system through an Ethernet connection using a secure shell client application. You use the RLM CLI to monitor and troubleshoot the storage system.



If you have a data center configuration where management traffic and data traffic are on separate networks, you can configure the RLM on the management network.

Note

In compliance with the GNU General Public License (GPL), Version 2, June 1991, a complete machine-readable copy of the source code for the relevant source code portions of the Remote LAN Module (RLM) Firmware that are covered by the GPL is available from http://now.netapp.com.

What the RLM does

The RLM provides the following remote management capabilities for the storage system. You use the RLM command line interface (CLI) commands to

- Remotely administer the storage system using the Data ONTAP CLI by using the RLM's system console redirection feature
- Remotely access the storage system and diagnose error conditions, even if the storage system has failed, by performing the following tasks:
  - View the storage system console messages, captured in the RLM's console log
  - View storage system events, captured in RLM's System Event Log
  - Initiate a storage system core dump
  - Power-cycle the storage system (or turn it on or off)
  - Reset the storage system
  - Reboot the storage system

  For more information, see "Troubleshooting the storage system with the RLM" on page 233.

The RLM extends AutoSupport capabilities by sending alerts or “down system” notification via an AutoSupport message when the storage system goes down, regardless of whether the storage system can send AutoSupport messages. For information about AutoSupport, see Chapter 8, “Using AutoSupport,” on page 163.

Configuring the RLM

Configuring the RLM for your storage system and network

Before using the RLM, you must configure it for your storage system and network. You can configure the RLM

- When setting up a new storage system with RLM already installed
- After setting up a new storage system with RLM already installed
- When you add an RLM to an existing compatible storage system

Prerequisites for configuring the RLM

Before you configure the RLM, you need to gather the following information:

- Network information
- AutoSupport information

Network information: You can configure the RLM using DHCP or static addressing.

- If you want to use DHCP addressing, bind the RLM's Media Access Control (MAC) address in the DHCP server. You can obtain the RLM's MAC address by using the Data ONTAP `rlm status` command at the storage system prompt. You can also obtain the RLM's MAC address by removing the RLM from the storage system and looking for the MAC address label on the module itself.
- If you want to use a static IP address, you need to gather the following information:
  - An available static IP address
  - Netmask of network
  - Gateway of network

AutoSupport information: RLM initiates and sends notifications, or event e-mail messages, using the settings of the following AutoSupport options:

- `autosupport.to` *`e-mail_addresses`*
- `autosupport.to { name | IP_address_of_outbound_SMTP }`

You must ensure the `autosupport.to` option has been set properly before configuring the RLM. You must enter the name or the IP address of the AutoSupport mail host when you configure the RLM. For more information about AutoSupport options, see "Configuring AutoSupport" on page 166.

Note

RLM does not rely on the storage system's `autosupport.support.transport` option to send "down filer" notifications. RLM uses the Simple Mail Transport Protocol (SMTP).

Configuring the RLM

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

- Initializing a storage system, with RLM pre-installed

  When a storage system is first installed, and power is applied, the storage system boots and displays a boot menu with several options, as described in Chapter 5, "Starting and Stopping the Storage System," on page 67. Select the appropriate option to continue the setup process. When the setup process is completed, the `rlm setup` command is automatically initiated. For more information about the entire setup process, see the *Software Setup Guide*.
- Running the Data ONTAP `setup` script

  This prompts you to configure the storage system, followed by initiating the `rlm setup` command.
- Running the Data ONTAP `rlm setup` command

When the RLM `setup` script is initiated, you are prompted to answer questions and enter network and mail host information. To configure the RLM, refer to the information you gathered in the section "Prerequisites for configuring the RLM" on page 207 and complete the following steps.

RLM configuration at power-up or with the setup command

To configure RLM when you first apply power to a storage system with RLM pre-installed or when you run the Data ONTAP `setup` command, complete the following steps.

Step	Action
1	Verify AutoSupport is configured properly.
2	At the storage system prompt, enter the following command: `toaster> setup`

Step	Action

	Example of using setup for with a static IP address:
	The following output is displayed:
	`The Remote LAN Module(RLM) provides remote management capabilities`
	`including console redirection, logging and power control.`
	`It also extends autosupport by sending`
	`additional system alerts. Your autosupport settings are used`
	`for sending these alerts via email over the RLM LAN interface.`
	`Would you like to configure the RLM? **y**`
	`Would you like to enable DHCP on the RLM LAN interface? **n**`
	`Please enter the IP address for the RLM []: **192.168.123.98**`
	`Please enter the netmask for the RLM []: **255.255.255.0**`
	`Please enter the IP address for the RLM gateway []: **192.168.123.1**`
	`The mail host is required by your system to send RLM`
	`alerts and local autosupport email.`
	`Please enter the name or IP address of the mail host []:`
	`Please enter the IP adcress for mailhost.xyz.com []:**10.57.160.114**`

	Example for using setup with a DHCP server:
	The following output is displayed:
	`The Remote LAN Module(RLM) provides remote management capabilities`
	`including console redirection, logging and power control.`
	`It also extends autosupport by sending`
	`additional system alerts. Your autosupport settings are used`
	`for sending these alerts via email over the RLM LAN interface.`
	`Would you like to configure the RLM? **y**`
	`Would you like to enable DHCP on the RLM LAN interface? **y**`
	`The mail host is required by your system to send RLM`
	`alerts and local autosupport email.`
	`Please enter the name or IP address of the mail host []:`
	`Please enter the IP adcress for mailhost.xyz.com []:**10.57.160.114**`

Step	Action

3	At the storage system prompt, enter the following command to verify the RLM's network configuration is correct or to display the MAC address of the RLM.
	`toaster> **rlm status**`
	Example: The following output is displayed if you used the static IP address in Step 2:
	`Remote LAN Manager`
	`Part Number: 111-00001`
	`Revision: 30`
	`Serial Number: 123456`
	`Firmware Version: 3.12`
	`Mgmt MAC Address: 00:AA:BB:CC:DD:EE`
	`Using DHCP: no`
	`IP Address: 192.168.123.98`
	`Netmask: 255.255.255.0`
	`Gateway: 192.168.123.1`

4	Verify that the RLM AutoSupport function is working properly by entering the following command. `toaster> **rlm test autosupport**` **Note** You must ensure the `autosupport.to` option has been set properly before issuing this command.

RLM configuration with the rlm setup command

To configure RLM when you first apply power to a storage system with RLM pre-installed or run the `setup` command, complete the following steps.

Step	Action
1	Verify AutoSupport is configured properly.
2	At the storage system prompt, enter the following command: `toaster> rlm setup`

	Example for using rlm setup with a static IP address:
	The following output is displayed:
	`The Remote LAN Module(RLM) provides remote management capabilities including console redirection, logging and power control. It also extends autosupport by sending additional system alerts. Your autosupport settings are used for sending these alerts via email over the RLM LAN interface.`
	`Would you like to configure the RLM? y`
	`Would you like to enable DHCP on the RLM LAN interface? n`
	`Please enter the IP address for the RLM []: 192.168.123.98`
	`Please enter the netmask for the RLM []: 255.255.255.0`
	`Please enter the IP address for the RLM gateway []: 192.168.123.1`

	Example for using rlm setup with a DHCP server:
	The following output is displayed:
	`The Remote LAN Module(RLM) provides remote management capabilities including console redirection, logging and power control. It also extends autosupport by sending additional system alerts. Your autosupport settings are used for sending these alerts via email over the RLM LAN interface.`
	`Would you like to configure the RLM? y`
	`Would you like to enable DHCP on the RLM LAN interface? y`

Step	Action

3	At the storage system prompt, enter the following command to verify the RLM's network configuration is correct or to display the MAC address of the RLM.
	`toaster> **rlm status**`
	Example: The following output is displayed if you used the static IP address in Step 2:
	`Remote LAN Manager`
	`Part Number: 111-00001`
	`Revision: 30`
	`Serial Number: 123456`
	`Firmware Version: 3.12`
	`Mgmt MAC Address: 00:AA:BB:CC:DD:EE`
	`Using DHCP: no`
	`IP Address: 192.168.123.98`
	`Netmask: 255.255.255.0`
	`Gateway: 192.168.123.1`

4	Verify that the RLM AutoSupport function is working properly by entering the following command. `toaster> **rlm test autosupport**` **Note** You must ensure the `autosupport.to` option has been set properly before issuing this command.

Managing the RLM with Data ONTAP

Data ONTAP commands and AutoSupport options

You manage the RLM from the storage system using the Data ONTAP `rlm` commands and changing the settings of AutoSupport settings that are used by the RLM.

Data ONTAP rlm commands

The following table describes the Data ONTAP `rlm` commands that you enter at the storage system prompt to manage the RLM. These commands are also described in the na_rlm man page.

Note

When you enter some of these commands, there might be a pause of a few seconds while the storage system queries the RLM. This is normal behavior.

Command	Description
`toaster> rlm help`	Displays the list of `rlm` commands available with the current release of Data ONTAP.
`toaster> rlm reboot`	Reboots the RLM and causes the RLM to perform a self-test. If your console connection is through the RLM, it will be disconnected.
`toaster> rlm setup`	Initiates the interactive RLM setup program. For more information, see "Configuring the RLM" on page 208.
`toaster> rlm status`	Displays the current status of the RLM. For more information, see Step 3 in "Configuring the RLM" on page 208.
`toaster> rlm test autosupport`	Sends a test e-mail to all recipients specified with the `autosupport.to` option.
`toaster> rlm update`	Updates the RLM firmware. For more information, see "The ways you can update the RLM firmware" on page 234.

AutoSupport options

You use the following options to manage AutoSupport event messages and alerts:

- `autosupport.to` *`e-mail_addresses`*
- `autosupport.mailhost {`*`name`* | *`IP_address_of_outbound_SMTP`*`}`

You use the following option to change the amount of information displayed when you use both the Data ONTAP and the RLM CLI commands to display status:

- `autosupport.content`

For more information, see the entry for `autosupport.content` in "AutoSupport options" on page 166 and the section "Contents of AutoSupport messages" on page 177.

Logging in to the RLM

Prerequisite for logging in to the RLM

Before logging in to the RLM, you must perform the following tasks:

- Install a secure shell client application that is appropriate for your administration host, such as SSH, OpenSSH for UNIX hosts, or PuTTY for Windows hosts.
- Ensure you have a user account and password with administrative privileges on the storage system. For information about storage system administrator accounts, see "About managing administrator access" on page 97.

  The RLM uses the same user credentials as the storage system. Changes to user account information on the storage system are updated to the RLM. User accounts cannot be created on the RLM.

Note

The RLM does not allow you to log in with the system administration account name of *root*. If you want to log in using the root account, use the name *naroot*. Then, if you are prompted for a password, type in the storage system root password.

Note

To provide additional security against unauthorized access to the storage system, create a password for the system administration account (root) for storage systems with an RLM installed.

For more information about how to create system administration accounts on storage system, see "About managing administrator access" on page 97.

For more information about configuring your storage system to use SSH with SecureAdmin, see Chapter 9, "Using SecureAdmin," on page 181 and the na_secureadmin(1) man page.

Logging in to the RLM from a UNIX host

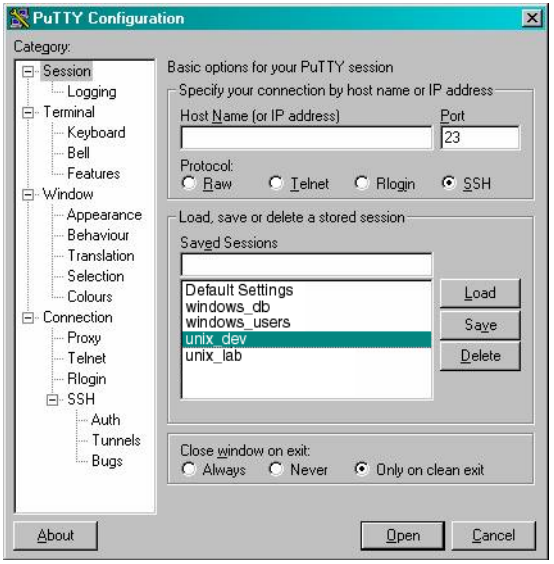
To log in to the RLM from a UNIX host, complete the following steps.

Step	Action
1	Ensure a secure shell application is installed on the UNIX host.

Step	Action
2	Enter the following command: `ra_host>` **`secure_shell_app username@RLM_IP_address`** **Example:** Logging in as root: `ra_host> ssh naroot@193.20.21.123` Logging in as root: `ra_host> ssh joe@193.20.21.123`

Logging in to the RLM from a Windows host

To log in to RLM from a Windows host, complete the following steps.

Step	Action
1	Download a secure shell application for Windows, such as PuTTY, and configure the IP address of the RLM.
2	Open a Windows session to log in to the RLM and make the appropriate selections. **Example:** 

Only one RLM session is allowed at a time

Only one administrator can be logged in to an active RLM CLI session at a time. If you have an RLM session open, you or another administrator with privileges to log in to the RLM can close your RLM session and open a new one. This is convenient if you logged into the RLM from one computer and forgot to close the session before moving to another computer, or if another administrator wants to take over the administration tasks from a different computer.

If a session is already initiated, you will see the following message:

```
User username has an active CLI session.
Would you like to disconnect that session, and start yours [y/n]?
```

If you enter `y`, the session owned by *username* is disconnected and your session is initiated. This action is recorded in the RLM's system event log.

If the username account for the storage system is password protected, you are prompted for the password. Enter the storage system password. When you are successfully connected, you see the RLM prompt.

The RLM prompt is created by adding "RLM" in front of the hostname of the storage system. For example, if your storage system is named toaster, the storage system prompt is `toaster>` and the prompt for the RLM session becomes `RLM toaster>`.

Managing the storage system with the RLM

Using the RLM command line interface

You manage the storage system with the RLM by using the RLM CLI, which has the same CLI features available in the Data ONTAP CLI:

- History
- Command-line editor
- Online command-line help

These CLI features that are common to Data ONTAP and the RLM are described in the section "Using the command line interface" on page 19. Note that RLM CLI commands are *not* documented in online command line manual (man) pages.

The RLM CLI provides two sets of commands, depending on the current user mode, or privilege level, which can be administrative (admin) or advanced. This functionality is identical to that of the privilege levels supported by the Data ONTAP CLI. For more information about privilege levels, see "Using Data ONTAP commands at different privilege levels" on page 22.

Using online help at the RLM CLI

You can display all of the available commands by entering the question mark (`?`) or `help`, as shown in the following example.

```
RLM toaster> ?
date
exit
events
help
priv
rlm
system
version
```

If a command has subcommands, they are displayed if you enter the command name after the `help` command, as shown in the following example:

```
RLM toaster> help events
events all
events info
events newest
events oldest
events search
```

For detailed online command help, you can enter the command followed by the question mark (`-?`) or the `-h` option. The basic command help is displayed if the command does not have subcommands. The following example shows the result of entering `-h` as an option for a command that has subcommands.

```
RLM toaster> events -h
events all - print all system events
events info - print system event log information
events newest - print newest system events
events oldest - print oldest system events
events search - search for and print system events
```

What you can do in RLM admin mode

In admin mode, you can use the RLM administrative commands to

- Connect to the storage system console (`system console`)
- Control the storage system power (`system power`)
- Display the following information:
  - Available commands (`help` or `?`)
  - Events that occur on the storage system (`events` *`subcommand`*)
  - Storage system console logs (`system log`)
  - Storage system power status (`system power status`)
  - Privilege level (`priv show`)
  - RLM status (`rlm status`)
  - RLM version (`version`)
  - Syntax usage for a specific command (`help` *`command`*)
- Dump the storage system core and reset the storage system (`system core`)
- Exit from the RLM CLI (`exit`)
- Reset the storage system with the firmware you specify (primary, backup, or current) (`system reset` *`firmware`*)
- Reboot the RLM (`rlm reboot`)
- Set the user mode privilege level (`priv set` *`level`*)
- Update RLM firmware (`rlm update` *`path`*)

Connecting to the storage system console

You use the `system console` command to log in to the storage system. When you are logged in to the storage system console, you use the Data ONTAP CLI to administer the storage system. You can enter commands at the storage system prompt or the CFE boot prompt. Use Ctrl-D to exit from the storage system console and return to the RLM CLI. For information about the procedure, see

"Using the RLM to remotely access the system console" on page 30. The following example shows the result of entering the `system console` command. Notice that the message "Type Ctrl-D to exit." appears and that you must press the Enter key to see the storage system prompt. Once the prompt appears, you can enter Data ONTAP commands. In the following example, the `vol status` command is entered, followed by Ctrl-D, which returns you to the RLM prompt.

```
RLM toaster> system console
Type Ctrl-D to exit.
```

(Press the Enter key to see the storage system prompt.)

```
toaster>
toaster> vol status
```

(information about all of the existing volumes is displayed)

```
toaster> Ctrl-D
RLM toaster>
```

Controlling storage system power

You use the `system power` command to turn the power on or off, or you can power-cycle the storage system, which automatically turns system power off and then back on. Turning the power off or power-cycling the storage system may cause an improper shutdown of the storage system (also called a dirty shutdown). The following example shows the result of using the s`ystem power cycle` command.

```
RLM toaster> system power cycle
This will cause a dirty shutdown of your appliance. Continue? [y/n]
```

If you press `y`, the storage system is turned off. Five seconds later, the storage system is turned back on, and the `CFE>` firmware prompt appears.

In the meantime, the RLM prompt awaits your next command.

```
RLM toaster>
```

Even when the storage system is powered off or power-cycled, the RLM remains operational.

Syntax summary for admin mode RLM CLI commands

The following table provides a quick reference of the command syntax for the RLM commands you can use in admin mode. For information on how to change the privilege level to admin mode, see "Setting the privilege level" on page 23.

Command syntax	Description
`RLM toaster> events {all | information | newest | oldest | search` *`string`*`}`	Displays storage system events stored by the RLM specified by the subcommand.
`RLM toaster> exit`	Exits from the RLM command line interface.
`RLM toaster> help [`*`command`*`]`	Displays a list of available commands, or if a command is specified, displays either the subcommands available for that command or its syntax usage.
`RLM toaster> priv set {admin | advanced | diag}`	Sets the privilege level to access the specified mode.
`RLM toaster> priv show`	Displays privilege level of the user mode.
`RLM toaster> rlm reboot`	Reboots the RLM. This action takes approximately one minute.
`RLM toaster> rlm sensors [-c]`	Displays RLM environmental sensor status. The `-c` option, which takes a few seconds to display, shows the current values, rather than cached values.
`RLM toaster> rlm status`	Displays RLM status. **Note** The Data ONTAP `sysconfig` command displays both the storage system and RLM status.
`RLM toaster> rlm update http://`*`path`*	Updates the RLM firmware.

Command syntax	Description

`RLM toaster> system console`
Result: The following prompt is displayed.
`toaster>`
To exit, enter Ctrl-D.
Result: The following prompt is displayed.
`RLM toaster>`
	Logs in to Data ONTAP at the current prompt. Use Ctrl-D to exit.

`RLM toaster> system core`	Dumps the storage system core and resets the storage system. This command is similar to pressing the Non-maskable Interrupt (NMI) button on a storage system. **Note** The RLM stays operational as long as input power to the storage system is not interrupted.
`RLM toaster> system power {on | off | cycle}`	Turns the storage system on or off, or performs a power cycle. When system power is turned off or power-cycled, standby power stays on. This may cause an improper shutdown of the storage system. When power-cycling, there is a 5-second pause before power is turned back on. **Note** The RLM stays operational as long as input power to the storage system is not interrupted.
`RLM toaster> system power status`	Displays status for each power supply, such as presence, input power, and output power.

Command syntax	Description
`RLM toaster> system reset {primary | backup | current}`	Resets the storage system using the specified firmware image. **Note** The RLM stays operational as long as input power to the storage system is not interrupted.
`RLM toaster> version`	Displays the RLM version information, including hardware and firmware information.

What you can display in RLM advanced mode

You can use the RLM advanced commands to display more information than is available in administrative mode:

- Debug log file (`rlm log debug`)
- RLM command history (`rlm log audit`)
- RLM environmental sensor status (`rlm sensors`)
- RLM messages file (`rlm log messages`)
- RLM statistics (`rlm status -v`)

Syntax summary for advanced mode commands

The following table provides a quick reference of the command syntax for the RLM commands you can use in advanced mode. For information on how to change the privilege level to advanced mode, see “Setting the privilege level” on page 23.

Command	Description
`RLM toaster*> rlm log audit`	Displays the RLM command history.
`RLM toaster*> rlm log debug`	Displays the RLM debug file.
`RLM toaster*> rlm log messages`	Dumps the RLM messages file.
`RLM toaster*> system sensors`	Displays a list of environmental sensors, their states, and their current values.

Displaying storage system and RLM information

About displaying information

The RLM provides several ways to display information about the storage system and the RLM itself. You can display real-time information using the commands in the admin or advanced mode, or you can display information stored in the RLM's System Event Log (SEL) or Console log. You can also view the information displayed in the additional AutoSupport messages generated by the RLM. Most of the information is stored in the SEL or in captured console messages.

Displaying real-time information

RLM CLI in admin mode: Using the RLM CLI commands in admin mode, you can view information about:

- The storage system power status (`system power status`)
- The status of the RLM (`rlm status`)
- The version of the RLM (`version`)

The following examples show how information is displayed using commands at the RLM admin mode prompt:

```
RLM toaster> system power status
Power supply1 status:
    Present: yes
    Turned on by Agent: yes
    Output power: yes
    Input power: yes
    Fault: no
Power supply 2 status:
    Present: yes
    Turned on by Agent: yes
    Output power: yes
    Input power: yes
    Fault: no

RLM toaster> rlm status
eth0  Link encap:Ethernet HWaddr 00:A0:98:01:9C:4B
      inet addr:10.41.42.73.231 Bcast:10.255.255.255
         Mask:255.255.255.0
      UP BROADCAST RUNNING MULTICAST   MTU:1500   Metric:1
      RX packets:8972  errors:0 dropped:0  overruns:0   frame:0
      TX  packets:72  errors:0  dropped:0  overruns:0  carrier:0
      collisions:0  txqueuelen:100
```

```
      RX bytes:622724  (608.1 kb)  TX bytes:8765 (8.5 kb)
      Interrupt:15

RLM toaster> version
serial#=123456
part#110-00030
rev#12
Agent revision: 12
Primary-
RLM_version=x.y (date)

Backup-
RLM_version=x.y (date)

Booted primary image
```

The RLM CLI in advanced mode: Using the RLM CLI commands in advanced mode, you can view information about:

- Internal RLM statistics (`rlm status -v`)
- The RLM environmental sensor (`rlm sensors`)

The following examples show how information is displayed using commands at the RLM admin mode prompt (note that the characters ". . ." indicate details have been omitted):

```
RLM toaster*> rlm status -v
eth0  Link encap:Ethernet HWaddr 00:A0:98:01:9C:4B
      inet addr:10.41.42.73.231 Bcast:10.255.255.255
         Mask:255.255.255.0
      UP BROADCAST RUNNING MULTICAST   MTU:1500   Metric:1
      RX packets:8972  errors:0 dropped:0  overruns:0   frame:0
      TX  packets:72  errors:0  dropped:0  overruns:0  carrier:0
      collisions:0  txqueuelen:100
      RX bytes:622724  (608.1 kb)  TX bytes:8765 (8.5 kb)
      Interrupt:15
packet reader daemon
---------------------------------
restarts                        1
port config errors              0
...
packet writer daemon
---------------------------------
restarts                        0
port config errors              0
...
console logger daemon
```

```
---------------------------------
logger restarts                  0
logger input packets             0
...
downbeat daemon
---------------------------------
Downbeat restarts                0
Downbeat packets                 0
...
upbeat daemon
---------------------------------
Upbeat restarts                  1
Upbeat packets                  93

ECC memory
---------------------------------
total corrections                0
totat uncorrectable errors       0
...
Watcher daemon
---------------------------------
watcher restarts                 0
agentd restarts                  0
...

RLM toaster*> rlm sensors
Sensor      Sensor  Sensor   Critical  Warning   Warning   Critical
Name        State   Reading  Low       Low       High      High
========    =====   ======   ======    ======    ======    =====
Temperature Normal  19C       N/A        0C         45C        60C
```

System Event Log

The RLM has a non-volatile memory buffer that stores up to 4,000 records of system events in a System Event Log (SEL). When the buffer is full, the oldest records are overwritten by the newest records.

The records contain the following data:

- Hardware events detected by the RLM—for example, system sensor status about power supplies, voltage, or other components
- Errors (generated by the storage system or the RLM) detected by the RLM—for example, a communication error, a fan failure, a memory or CPU error, or a “boot image not found” message
- Critical software events sent to the RLM by the storage system—for example, a system panic, a communication failure, an unexpected CFE

prompt, a CFE boot failure, or a user-triggered “down system” as a result of issuing the `system reset` or `system power cycle` command

Note

The SEL uses the RLM’s clock to time-stamp events. RLM begins synchronizing its clock with the system clock as soon as the storage system boots up. However, synchronizing takes a few seconds. If events occur during these few seconds, they are time-stamped January 1, 1970.

Recent records from the SEL are attached to the AutoSupport messages sent by the RLM.

You access the SEL by using the RLM `events` command, as shown in the following examples.

```
RLM toaster> events info
SEL version: 4
Maximum number of records: 4000

Number of records: 180
Record size: 64 Bytes
Oldest record’s id: 1
Latest records’s id: 180

RLM toaster> events oldest 1
Record 1: Wed Mar 30 20:17:38 2005 [RLM CLI.notice]: “log in”

RLM toaster> events all
```

(displays all events)

```
RLM toaster> events newest 2
Record 179: Wed Mar 30 10:32:49 2005 [Agent Event.warning]: FIFO
    0x001A - Agent XYZ, PS2_PRSNT easserted.
Record 180: Wed Mar 30 10:32:50 2005 [Agent Event.warning]: FIFO
    0x8017 - Agent XYZ, FAN_ALERT2 easserted

RLM toaster> events search WD
Record 5: Tue Mar 29 07:39:40 2005 [Agent Event.warning]: FIFO
0x8FFF - Agent XYZ, L1_WD_TIMEOUT asserted.
Record 6: Tue Mar 29 07:39:42 2005 [Agent Event.critical]: FIFO
0x8FFE - Agent XYZ, L2_WD_TIMEOUT asserted
```

Console log

The RLM monitors the storage system console regardless of whether any administrators are logged in or connected to the console. When storage system messages are sent to the console, the RLM stores them in the console log, which resides in a 96-KB buffer in its main memory. The console log can store approximately 2,000 lines of system console messages. When the buffer is full, the oldest messages are overwritten by the newest messages.

The console log persists as long as the RLM has power from either of the storage system's power supplies. Since the RLM operates with standby power, it remains available even when the storage system is power-cycled or turned off.

If the `autosupport.content` option is set to `complete`, and a "down filer," a system hang, or a reboot loop condition occurs, the console logs are attached to the AutoSupport messages sent by the RLM.

You display the contents of the console log with the `system log` command, as shown in the following example.

```
RLM toaster> system log

toaster> Wed Mar 30 10:32:49 GMT
    [monitor.fan.critical:EMERGENCY] : SYS 2 FAN 2 is spinning
    below the tolerable speed. Replace it immediately to avoid
    overheating.
toaster> Wed Mar 30 10:33:10 GMT
    [monitor.fan.critical:EMERGENCY] : current speed is 82 RPM,
    critical low is 400 RPM
```

AutoSupport messages

For storage systems with RLM, there are two additional types of AutoSupport messages:

- RLM-generated AutoSupport messages about the storage system
- Storage system-generated AutoSupport messages about the RLM

For information about how to interpret AutoSupport messages, see "Contents of AutoSupport messages" on page 177.

RLM-generated AutoSupport messages: These messages include the following information:

- In the subject line—A system notification from the RLM of the storage system, listing the system condition or event that caused the AutoSupport message, and the log level (for information about log level, see "Event and log-level descriptions" on page 176)
- In the message body—The RLM configuration and version information, the storage system ID, serial number, model and host name

- In the gzipped attachments—the System Event Logs, the system sensor state as determined by the RLM, and console logs. (Console logs can be omitted by setting the `autosupport.content` option to `minimal`)

Typical RLM-generated AutoSupport messages occur in the following conditions:

- The storage system reboots unexpectedly.
- The storage system stops communicating with the RLM.
- A Watchdog reset occurs.
- The storage system is power-cycled.
- Firmware POST errors occur.
- A user-initiated AutoSupport message occurs.

Storage system-generated AutoSupport messages: include the following information:

- In the subject line—A system notification from the name of the storage system with the RLM, a description of the RLM condition or event that caused the AutoSupport message, and the log level (for information about log level, see "Event and log-level descriptions" on page 176)
- In the message body—A time stamp, the system software version and storage system ID, host name, and output from the `sysconfig -a` command
- In the gzipped attachments—messages from EMS, rc, exports, hosts, resolv_conf, nsswitch_conf, and cm_stats

Typical storage system AutoSupport messages about the RLM occur in the following conditions:

- The RLM stops communication with the storage system.
- The RLM software fails.
- The RLM hardware fails.

Comparing Data ONTAP and RLM commands

The following table shows the Data ONTAP commands that are used to manage the RLM and the RLM commands that are used to manage the storage system.

Action	Data ONTAP Command or Procedure	RLM Command
Set up RLM in a new storage system	Turn on the new storage system. During bootup, press Ctrl-C. From the menu, select `4` or `4a` to initialize disks and set up a root volume. **Result:** After the Data ONTAP `setup` script is completed, the RLM `rlm setup` script is initiated.	
Reconfigure an RLM in an existing storage system	`toaster>` **`setup`** **Result:** After the Data ONTAP `setup` script is completed, the `rlm setup` script is initiated.	
Initiate the `rlm setup` script to configure RLM	`toaster>` **`rlm setup`**	
Test the RLM's AutoSupport setting	`toaster>` **`rlm test autosupport`**	
Display Data ONTAP `rlm` commands	`toaster>` **`rlm help`** `help setup test update reboot status`	
Log in to the RLM		`ra_host>` **`ssh naroot@RLM_IP_addr`** `ra_host>` **`ssh user@RLM_IP_addr`** `RLM toaster>`

Action	Data ONTAP Command or Procedure	RLM Command

Display RLM CLI commands		`RLM toaster>` **`?`**
		`date`
		`exit`
		`events`
		`help`
		`priv`
		`rlm`
		`system`

Display the twenty most recent events logged by RLM		`RLM toaster>` **`events newest 20`**
Display a summary of information about the records in the events log		`RLM toaster>` **`events info`**

Display the RLM configuration	`toaster> rlm status`
	`toaster*> sysconfig -v`
	Requires advanced mode.
		`RLM toaster>` **`rlm status`**

Display statistics gathered by RLM	`toaster*> rlm status -v` Requires advanced mode.	`RLM toaster*>` **`rlm status -v`** Requires advanced mode.
Display the system hardware sensor list		`RLM toaster*>` **`system sensors`** Requires advanced mode.
Log in to the system to manage storage system resources		`RLM toaster>` **`system console`** toaster> (Use Ctrl-D to exit).
Dump the system core and reset the storage system		`RLM toaster>` **`system core`**
Reset the RLM	`toaster>` **`rlm reboot`**	`RLM toaster>` **`rlm reboot`**

Action	Data ONTAP Command or Procedure	RLM Command
Update the RLM firmware	toaster> **software install http://path/RLM_FW.zip -f** toaster> **rlm update**	RLM toaster> **rlm update http://*path*/RLM_FW.tar.gz**

Troubleshooting the storage system with the RLM

About troubleshooting

When you become aware of a problem with the storage system, you can use the RLM to display information about the problem, dump a system core, and reboot the storage system, even if the storage system's firmware is corrupted. Use the following table as a guideline for troubleshooting a storage system.

Note

If you want the RLM to send you AutoSupport messages to keep you informed of the status of both the storage system and the RLM, properly configure the AutoSupport feature.

If this condition occurs...	And you want to...	Enter this command at the RLM CLI prompt...
The storage system is not responding properly	Access the storage system console	`system console`
You receive an AutoSupport message for an event that is occurring or has occurred, such as the failure of a hardware component or a storage system that has panicked and is down.	Display what has occurred at the storage system console	`system log`
	Display all events, starting with most recent	`events all`
	Display a specific number of recent events	`events newest` *`number`*
	Search for specific events in the SEL	`events search` *`string`*
The storage system is hanging	Dump the system core and reboot the storage system	`system core`
	Power-cycle the storage system	`system power cycle`
The storage system firmware is corrupted	Boot using a backup copy of the storage system firmware	`system reset backup`

Updating the RLM firmware

The ways you can update the RLM firmware

You can download and update the RLM firmware from the Data ONTAP command line interface (CLI) or the RLM CLI.

- If you want to use the Data ONTAP CLI, use the .zip file.

  Note

  Before you can upgrade the RLM firmware from the Data ONTAP CLI, you must run the `software install` command to download the appropriate software. For information about the `software` command and its subcommands, see the *Upgrade Guide* and the na_software (1) man page.

- If you want to use the RLM CLI, use the .tar.gz.file.

When you download the RLM firmware from the NOW site at http://now.netapp.com, you can save the file on any web server that is accessible from the storage system.

Note

After the firmware is updated, you are prompted to reboot the RLM. If your console connection to the storage system is being redirected through the RLM, you will lose your connection. Wait for approximately one minute for the RLM to reboot and re-establish your console connection through the RLM.

Prerequisites

You must have the following items to download and update the firmware:

- Access to a web server on a network accessible to your storage system
- The name and IP address of the web server
- Access to the storage system's serial console

Using the Data ONTAP CLI: To update the RLM firmware at the storage system prompt, complete the following steps:

Step	Action
1	Go to Firmware Instructions for the Remote LAN Module at the NOW site at http://now.netapp.com/.
2	Click on the RLM_FM.zip link to download the file from the web site to the web server on a network that is accessible to your storage system.

3	Enter the following command:
	`toaster> **software install http://web_server/RLM_FW.zip -f**`
	web_server is the name of the web server on a network accessible to your storage system.
	Result: The following messages are displayed.
	`software: copying to /etc/software/RLM_FM.zip`
	`software: 100% file read from location.`
	`software: /etc/software/RLM_FM.zip has been copied.`
	`software: installing software, this could take a few minutes...`

4	Enter the following command:
	`toaster> **rlm update**`
	Result: The following messages are displayed.
	`Updating the RLM firmware.`
	`DO NOT reset this system during this process.`
	`New RLM version : x.x.x`
	`Sending file to RLM. . .`
	`Current RLM version : x.x.x`
	`Installing package on RLM. . .`
	`RLM: Firmware updated successfully!`

Step	Action
5	When the system prompts you to update the RLM, enter `y` to continue. **Result:** The RLM is updated and you are prompted to reboot the RLM. Wait 60 seconds to allow the RLM to reboot. **Note** Assuming that your console connection is not through the RLM, entering `y` to reboot the RLM, will not cause you to lose your console connection to the storage system.

Using the RLM CLI: To update the RLM firmware at the RLM prompt, complete the following steps.

Step	Action
1	Go to Firmware Instructions for the Remote LAN Module at the NOW site at http://now.netapp.com/.
2	Click on the RLM_FM.tar.zip link to download the file from the web site to the web server on a network that is accessible to your storage system.
3	Log in to the RLM by entering the following command at the administration host. `admin_host> ssh username@RLM_IP_address`

Step	Action

4	Enter the following command:
	`RLM toaster> rlm update http://web_server_addr/RLM.FW.tar.gz`
	web_server_addr is the IP address of the web server on a network accessible to your storage system.
	Result: The following messages are displayed.
	`Downloading package. . .`
	`Uncompressing package. . .`
	`Unarchiving package. . .`
	`Installing package. . .`
	`Reprocessing configuration state. . .`
	`Cleaning up. ..`
	`RLM FW Update Successful.`

5	When you are prompted to reboot the RLM, enter the following command: `RLM toaster> rlm reboot` **Note** If your console connection is through the RLM, you will lose your console connection to the storage system. Wait for approximately one minute for the RLM to reboot and re-establish your console session through the RLM.

For the latest information

For the latest information about updating RLM firmware, see the NOW site at http://now.netapp.com.

Troubleshooting RLM problems

Communication problems

If there is a communication failure between the storage system and the RLM, search for EMS events titled:

`[rlm.orftp.failed:warning]: RLM communication error, (`*`reason`*`)`

Configuration problems

If you are having difficulty configuring the RLM, view detailed configuration information and assign a new IP address by completing the following steps.

Step	Action
1	Verify the RLM is online and the IP configuration is correct by entering the following command: `toaster*> rlm status`
2	If the RLM is configured using DHCP, reconfigure the RLM using a static IP address by entering the following command: `toaster*> rlm setup`

Connection problems

If you are having difficulty connecting to the RLM, check the connection by completing the following steps.

Step	Action
1	Verify you are using a secure shell client to connect to the RLM. For information on secure shell client applications, see "Prerequisite for logging in to the RLM" on page 215.
2	From the storage system, verify the RLM is online and the IP configuration is correct by entering the following command: `toaster> rlm status`
3	From the administration host, test the network connection for the RLM by entering the following command: `ra_host> ping rlm_IP_address` If the ping fails: ◆ Verify the RLM network port on the back of the storage system is cabled and active. For more information, see the Installation and Setup Instructions for your storage system. ◆ Verify the RLM has a valid IP address. (At the storage system, enter the `rlm setup` command to use the DHCP server or assign a valid IP address.). ◆ Verify the administration host has a route to the RLM.
4	From the storage system, reboot the RLM by entering the following command: `toaster> rlm reboot` **Note** It takes approximately one minute for the RLM to reboot.
5	If the RLM does not reboot, repeat Steps 2 through 4. If the RLM still does not reboot, contact technical support for assistance.

RLM failures

An RLM failure can occur for one of the following reasons:

- A communication link between the RLM and the storage system is broken.
- A hardware failure has occurred on the RLM.

If the RLM fails, search for EMS events titled:

`[rlm.heartbeat.stopped:warning]: Have not received a Heartbeat from the Remote LAN Module in the last` *n* `seconds, (`*reason*`)`

To troubleshoot an RLM failure, complete the following steps.

Step	Action

1	Run diagnostics by entering the following command from the CFE prompt:
	`CFE>` **`boot_diags`**
	Result: The diagnostics main menu appears.
	`Copyright (c) 1992-2005 Network Appliance, Inc.`
	`Diagnostic Monitor v-4. 4x67`
	`----------`
	`all Run all system diagnostics`
	`mb FAS3020 motherboard diagnostic`
	`mem main memory diagnostic`
	`agent agent & rlm diagnostic`
	`cf-card CompactFlash controller diagnostic`
	`fcal FCAL controller diagnostic`
	`stress System wide stress diagnostic`
	`Commands:`
	`Config (print a list of configured PCI devices)`
	`Default (restore all options to default settings)`
	`Exit (exit diagnostics and return to firmware OK prompt)`

Step	Action

2	From the main menu, enter the following option:
	`Enter Diag, Command or Option:` **`agent`**
	Result: The following RLM diagnostic menu appears.
	`Agent Diagnostic`
	`-------------------------`
	` 1: Comprehensive test`
	` 2: Appl-Agent interface test`
	` 3: Appl PS On-Off test  70: Show Agent ring buffer info`
	` 4: RLM Memory test  71: Show RLM info`
	` 5: RLM Sensor test  72: Show Restart reason`
	` 6: RLM-Agent interface test`
	` 7: RLM IRQ test`
	` 8: RLM NMI test  91: Enable/disable looping`
	`  92: Stop/continue on error`
	`11: RLM PS On-OFf test  93: Extended/Normal test`
	`  99: Exit`
	`Select test or feature by number [0]:`

3	From the RLM diagnostic prompt, enter test number 1 as shown: `Select test or feature by number [0]:` **`1`** **Note** It takes approximately ten minutes to complete this test. **Result:** This initiates a comprehensive test that includes running tests 2 through 8, and 11. The results of each test are displayed.
4	Based on the results of Step 3, diagnose the problem. If the problem persists, reseat the RLM and repeat Steps 1 to 4. If the problem still persists, replace the RLM.

RLM firmware update problems

A firmware update failure can occur for one of the following reasons:

- The firmware image is incorrect or corrupted.
- A communication error occurred while sending firmware to the RLM.
- The update failed when you attempted to install the new firmware at the RLM.
- The storage system was reset during the update.
- There was a power loss during update.

To troubleshoot a firmware failure, search for EMS events titled: `[rlm.orftp.failed:warning]: RLM firmware update failed, (`*reason*`)`

Example: `[rlm.orftp.failed:warning]: RLM firmware update failed: ORFTP couldn't send info.symlinks to RLM.`

This error message indicates that the firmware update failure is due to a communication error.

Troubleshooting with the Data ONTAP CLI. To troubleshoot updating the firmware using the Data ONTAP CLI, complete the following steps.

Step	Action
1	Verify the RLM is online by entering the following command: `toaster> rlm status`
2	Update the firmware by entering the commands described at "Using the Data ONTAP CLI" on page 235. **Note** Verify you are using the correct filename (*filename*.zip) of the RLM firmware.
3	Reboot RLM by entering the following command: `toaster> rlm reboot` **Note** It takes approximately one minute for the RLM to reboot.
4	If the RLM does not reboot, repeat Steps 1 through 3. If the RLM still does not reboot, contact technical support for assistance.

Troubleshooting with the RLM CLI: To troubleshoot updating the firmware using the RLM CLI, complete the following steps.

Step	Action
1	Verify the RLM is online by entering the following command: `toaster> rlm status`
2	From a browser, access the RLM firmware file on your Web server. **Note** Verify you are using the correct filename (*filename*.tar.gz) of the RLM firmware.
3	Update the firmware by entering the following command. `RLM toaster> rlm update http://path_hostname/RLM.FW.tar.gz` If this command fails, replace the *path_hostname* with the correct IP address.
4	Reboot the RLM by entering the following command: `toaster> rlm reboot`

System Information and Performance 11

About this chapter

This chapter describes commands that give you information about the storage system and how it is performing. It also describes some tips on improving storage system performance.

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:

- "Storage system configuration information" on page 246
- "Storage information" on page 249
- "Aggregate information" on page 251
- "Volume information" on page 253
- "File statistics for volumes" on page 255
- "Environmental status information" on page 262
- "Fibre Channel information" on page 265
- "Getting storage system information using the stats command" on page 266
- "Getting system information using perfmon" on page 277
- "Improving storage system performance" on page 278

Storage system configuration information

Version information

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

Hardware configuration

The `sysconfig` command displays information about the storage 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 storage 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 storage 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

Note

You can also obtain the information displayed by `sysconfig -r` from SNMP, using the custom Management Information Base (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 storage system was booted with the autoload setting of the tape library off.

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

Comprehensive storage system information

The `sysconfig -A` command displays storage 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 storage system information using the stats command" on page 266.

Storage information

Command for displaying storage information

You use the `storage show` command to display information about storage components for a storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 aggregates, see the chapter on aggregates in the *Storage Administration Guide*. 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 RAID4 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 storage system information using the stats command" on page 266.

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 storage system is displayed.

For more information about volumes, see the chapter on volumes in the *Storage Management Guide*. 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 volume 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 (through the `aggr options` or `vol options` commands).

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 storage system information using the stats command" on page 266.

File statistics for volumes

About file statistics

The `filestats` command displays a summary of file statistics within a volume on a storage system by reading file information from a Snapshot copy that you specify. File statistics help you determine when to schedule Snapshot copies by enabling you to see when most file activity takes place on a volume. Using the `filestats` command also helps you determine Snapshot copy disk consumption.

Detailed information

This section contains the following topics:

- "About the filestats command" on page 256
- "Options for the filestats command" on page 259

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 copy 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 copy.

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 259.

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 copy.

Example with no options specified

The following example shows sample output from the `filestats` command, without any options, for the hourly.1 Snapshot copy 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 storage 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 Snapshot copies 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 Snapshot copies, divide 40 by 3 to determine the number of daily Snapshot copies you can have before exceeding your space limit. In this example, 13 daily Snapshot copies 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 Snapshot copies.

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 Snapshot copies during these hours and fewer throughout the rest of the day. Scheduling more Snapshot copies before or during increased file activity decreases the time between file changes and Snapshot copies.

For information about managing Snapshot copy, 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 257.

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).

Note

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 257.

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 storage 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 storage system environment:

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

You can query information about the following items:

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

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

All environment information

The `environment status` command with no options displays all storage 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 storage 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 storage system information using the stats command" on page 266.

Getting storage 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 predefined data collection tools in Data ONTAP called counters. These counters provide you with information about your storage 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 storage 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 storage 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 storage 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 266.
- 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?

  Some preset files are provided with Data ONTAP. You can also create your own.

  For more information, see "About the stats preset files" on page 275.
- 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 274.

- 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 271.

  - 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 272.

  - 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 273.

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 271.

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.

Note

Each instance of a `stats` command consumes a small amount of system resources. If you start a large number of `stats` commands in background mode, you could affect overall storage system performance. To avoid this issue, make sure that you stop all background `stats` commands that you start.

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 271.

Step	Action
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.

If you are displaying multiple objects that have different counters, the column format may be difficult to read. In this case, use the row format.

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

Data ONTAP 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.

Getting system information using perfmon

About using perfmon to access system performance statistics

The `perfmon` performance monitoring tool is integrated with the Microsoft Windows operating system. If you use storage systems in a Windows environment, you can use `perfmon` to access many of the counters and objects available through the Data ONTAP `stats` command.

Using perfmon to access system performance statistics

To use `perfmon` to access storage system performance statistics, specify the name or IP address of the storage system as the counter source. The lists of performance objects and counters then reflect the objects and counters available from Data ONTAP.

Note

The default sample rate for `perfmon` is once every second. Depending on which counters you choose to monitor, that sample rate could cause a small performance degradation on the storage system. If you want to use `perfmon` to monitor storage system performance, you are advised to change the sample rate to once every ten seconds. You can do this using the System Monitor Properties.

Improving storage 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 storage 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 storage 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, you are advised to support 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 storage 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.

Attention

If you are not sure whether your storage 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 three, setting minimal read-ahead to On might improve performance. By default, the storage 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 storage system does very aggressive read-ahead.

Adding disks to a disk-bound aggregate

If you have a single traditional volume or single-aggregate storage 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 (two disks for RAID4 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 information about 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 Snapshot copies. 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 storage systems

Oplocks allow CIFS clients to cache more data locally, reducing traffic to the storage 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.

Attention

You are strongly advised to call 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 storage system and the CIFS client at one time.

Attention

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 storage 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 storage 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 storage 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:

- "Storage system panics" on page 284
- "Error messages" on page 285
- "Other troubleshooting resources" on page 287

Storage system panics

About storage system panics

If your storage 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 storage system reboots.

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

What to do if your storage system panics

If your storage system panics, following these steps can help 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 web 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 web site and upgrade to resolve the issue.
5	Otherwise, call technical support.

Note

If you have AutoSupport enabled, AutoSupport automatically alerts technical support 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 storage 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 145.

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 web site.

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

Step	Action
1	Go the web 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 storage 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.

Using RLM to troubleshoot

If your storage system supports the Remote LAN Module (RLM), you can use it to troubleshoot the system even if you are not in the same location as the system. You can use the RLM to view system console messages, view system events, dump the system core, and issue commands to power-cycle, reset, or reboot the system. For more information, see “Troubleshooting the storage system with the RLM” on page 233.

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 an 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 Manage ONTAP.

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 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 163.

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 storage system.

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

cluster interconnect Cables and adapters with which the two storage 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 storage systems in the cluster through the `cf` command.

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

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

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

copy-on-write The technique for creating Snapshot copies without consuming excess disk space.

degraded mode The operating mode of a storage system when a disk is missing from a RAID4 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 storage 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 storage system.

emulated storage system A software copy of a failed storage system that is hosted by its takeover storage system. The emulated storage system appears to users and administrators to be a functional version of the failed storage system. For example, it has the same name as the failed storage 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 storage system expansion slot. See adapter card.

expansion slot The slots on the storage system board into which you insert expansion cards.

failed storage system A physical storage system that has ceased operating. In a clustered configuration, it remains the failed storage 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 storage system to the failed storage system, resulting in a return to normal operation; the reverse of takeover.

group A group of users defined in the storage system's /etc/group file.

heartbeat A repeating signal transmitted from one storage system to the other that indicates that the storage system is in operation. Heartbeat information is also stored on disk.

hot spare disk A disk installed in the storage 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 storage 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 storage system and in a UNIX file system.

interrupt switch A switch on some storage 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 storage system The storage 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 storage system that is used to store the cluster state information of a storage system. If that system stops operating, the takeover system uses the information in the mailbox disks in constructing a virtual storage system. Mailbox disks are also used as file system disks.

maintenance mode An option when booting a storage 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 storage system so that it appears as multiple storage systems on the network.

NDMP Network Data Management Protocol. A protocol that allows storage 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 storage system when there is no takeover in the cluster.

NVRAM cache Nonvolatile RAM in a storage system, used for logging incoming write data and NFS requests. Improves system performance and prevents loss of data in case of a storage system or power failure.

NVRAM card An adapter card that contains the storage system's NVRAM cache.

NVRAM mirror A synchronously updated copy of the contents of the storage system NVRAM (nonvolatile random access memory) contents kept on the partner storage system.

panic A serious error condition causing the storage 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 RAID4 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 storage system, the other storage system in a cluster.

partner mode The method you use to communicate through the command-line interface with a virtual storage system during a takeover.

POST Power-on self-tests. The tests run by a storage 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. storage 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 storage system models.

serial console An ASCII or ANSI terminal attached to a storage system's serial port. Used to monitor and manage storage system operations.

share A directory or directory structure on the storage 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 copy An online, read-only copy of an entire file system that protects against accidental deletions or modifications of files without duplicating file contents. Snapshot copy enables users to restore files and to back up the storage system to tape while the storage system is in use.

system board A printed circuit board that contains a storage system's CPU, expansion bus slots, and system memory.

takeover The emulation of the failed node identity by the takeover node in a cluster; the opposite of *giveback*.

takeover storage system A storage system that remains in operation after the other storage system stops working and that hosts a virtual storage system that manages access to the failed node disk shelves and network connections. The takeover node maintains its own identity and the virtual node maintains the failed node identity.

takeover mode The method you use to interact with a storage system while it has taken over its partner. The console prompt indicates when the storage 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 storage 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 unit A virtual storage system you create using MultiStore, which enables you to partition the storage and network resources of a single storage system so that it appears as multiple storage 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 storage 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 10, 59
/etc/hosts.equiv file 47
/etc/log/auditlog file 149
/etc/messages file 59, 145
/etc/rc file 152, 155
/etc/syslog.conf file
    configuring message logging in 148
    file format and parameters of 146
/etc/usermap.cfg file, character coding of 59
/home file, contents of 58
/vol/vol0, root volume 55

Numerics

3DES, for SecureAdmin 183

A

access to storage system *See* storage system access
ACL, definition of 289
adapter card, definition of 289
address resolution, definition of 289
administration hosts
    adding 47
    definition of 45, 289
    deletion of 48
    reasons for using 45
    removing 48
    use of 45
    when you can add 45
    where they are specified 46
administrative level commands 22
administrator access, managing 97
administrator accounts
    changing the password of (passwd) 126
    reasons for creating 97
advanced level commands 22
aggregate Snapshot copy management 132
aggregates
    aggr copy command 10
    aggr status command, description of 251
    aggregate state, displaying (aggr status) 251
    definition of 289
    disk statistics, displaying (aggr status) 251
    performance improvements for disk-bound aggregates 279
    root option 66
API, definition of 289
ATM, definition of 289
auditlog file 97, 149
authentication
    public key-based 193
    with SSH 183
    with SSL 186
authentication, definition of 290
AutoSupport
    about 164
    commands used 178
    configuring 170
    contents of email 177
    default transport protocol 164
    defined 164
    events that trigger e-mail 176
    mail host support for 165
    options 166
    options AutoSupport.option (configures AutoSupport) 170
    reboots and 144
    requirements for 165
    technical support and 164
    testing 170
    testing (options AutoSupport.doit) 170
    transport protocols 165
    troubleshooting 172
    when system reboots 144

B

banner message for Telnet sessions 35
big-endian, definition of 290
board (storage system hardware) 4
boot options
    FAS200 series storage systems 75

for F87 systems 73
booting
F825, F880, FAS series systems 85
F87 systems 73
FAS200 series systems 75
from command prompt 68
from diskette 70, 71, 78, 85, 87
from firmware prompt 68, 76
options for the F87 storage system 73, 75
reasons to 70, 85
remotely 79
browsers, improving security through 186
Bugs Online 287

C

capabilities
assigning to roles 109
assignment to users 98
definition of 98
list of supported types 109
modifying others' 104
types of 109
cards, expansion, displaying information about 246
certificate-authority signed certificates 197
certificates
domain names and 200
generating 197
installing 197, 199, 201
regenerating 201
testing 200
types of 197
used by SSL protocol 197
change privileges, file ownership 52
character coding for configuration files 59
checksums, displaying information 251
CIFS
accessing /etc directory 61
accessing home directory 62
administrator accounts in 97
client, requirements to manage storage system 46
client,increasing performance for 280
editing configuration files using 63
CIFS share, definition of 295
CIFS, definition of 290
CLI, definition of 290
client decryption 184
client, definition of 290
clients
CIFS, requirements 46
editing configuration file from 63
NFS, requirements 46
platforms supported by FilerView 40
SecureAdmin supported 183
cluster interconnect, definition of 290
cluster monitor, definition of 290
cluster, definition of 290
command line editing
with editing mechanism 19
with history mechanism 19
commands
administrative level 22
advanced level 22
AutoSupport.option (sets AutoSupport options) 170
boot (boots the storage system) 68
date (sets system date and time) 138
halt (halts the storage system) 93, 94
license 136
options AutoSupport.doit (tests AutoSupport) 170
options wafl.root_only_chown (sets file ownerships changes) 52
passwd (changes administrative user password) 126
passwd (changes storage system system password) 125
priv 22
privilege levels 22
reboot (reboots the storage system) 91
rsh command list 38
savecore, what it does 143
See also rsh commands, statistics commands, stats command, status commands, SSH (Secure Shell) commands, SSL (secure sockets layer) commands.
stats 266
useradmin 97
community, definition of 290

CompactFlash cards
booting an F87 storage system 73
booting F800 systems 85
booting FAS900 systems 85
booting the FAS200 series storage systems 75
checking the Data ONTAP version of 81, 88
description of 4
recovering from corrupted image 81
configuration
display, using sysconfig 246
message logging 145
of AutoSupport (options AutoSupport.option) 170
configuration files
/etc 58
about backup and cloning 156
accessing 17
backing up 156
cloning 157
comparing backups 158
editing from CIFS client 63
editing from NFS client-setup 63
restoring 157
within /etc directory 59
configuring
HTTP services 88
TFTP services 89
console, definition of 291
continuous media scrub, definition of 291
copy-on-write, definition of 291
core files 143
criticaltime (UPS option) 159

D

data access management 10
data migration management 10
Data ONTAP, check version of 81
data organization management 9
data protection 10
data storage management 9
DataFabric Manager 2
date, setting storage system time and 138
decryption, between client and storage system 184
default directories 58
default root volume 55
degraded mode, definition of 291
deletion of administration hosts 48
device carrier 3
directories, default permissions 58
disk ID number, definition of 291
disk sanitization, definition of 291
disk shelf, definition of 291
disks, displaying statistical information for 246
domain names, changing storage system 200
domainusers
definition of 97
deleting 119
granting access to 102
listing 114
DSA key pair, for SSH 193

E

editing mechanism, editing command line with 19
emulated storage system, definition of 291
encryption
with SSH 184
with SSL 186
encryption algoritims supported by SecureAdmin 183
environmental adapter (storage system main unit) 5
error message logging, about 285
expansion slots 6

F

F800, booting with CompactFlash 85
F87 storage system, booting 73
FAS200 series systems
booting 75
preparing for Netboot startup 82
starting through netboot option 83
FAS900, booting with CompactFlash cards 85
file ownership change privileges 52
FilerView
accessing storage system through 17, 40
accessing Syslog Translator through 286
description 40
Help system defined 44
supported by client platforms 40

using the interface 42
Web browser requirements 40
files, configuration 58
filestats command
about 256
options for 259

G

generating certificates 197
groups
assigning roles to 106
assigning users to 100
definition of 98
deleting 119
listing 114
naming requirements 98
predefined 105
renaming 107
Windows special 99

H

history mechanism, editing command line with 19
HMAC, for SecureAdmin 183
host keys
changing the size of 190
determining sizes 188
setting 188
uses of 183
using with SSH 188
where stored 188
host support for AutoSupport, mail 165
hosts
definition of 45
deletion of administration 48
use of administration 45
when you can add 45
HTTP access to log files 62
HTTP services, configuring 88

I

installing certificates 197, 199
interface, use of FilerView 42

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 183
used by SSH protocol 183
Knowledgebase Solutions 287

L

LCD, on storage system chassis 5
level commands
administrative 22
advanced 22
licenses 136, 137
log files, accessing using HTTP 62

M

mail host support for AutoSupport 165
Manage ONTAP Developer SDK software 15, 16, 17
man-in-the-middle warning 191
manual (man) pages 21
memory (storage system main unit) 4
message files, accessing using HTTP 62
message logging, configuring 145
MetroCluster 11
mount privileges, controlling of (options nfs.mount_rootonly) 51
multiprotocol file and block sharing 8

N

N5200 and N5500 systems, booting 91
naming requirements for useradmin command 98
NDMP 11
Netboot startup
preparing FAS200 series storage systems 82
server requirements 88
starting FAS200 series systems 83
using storage systems as Netboot servers 88
Network file service 8

NFS client
access to /etc directory 61
access to /etc/directory 62
requirements to manage storage system 46
non-local users, granting access to 102
NVFAIL 11
NVRAM
description of 4
halt command to save data to disk 93

O

obsolete domain names, and SSL 200
OpenSSH
for SecureAdmin 183
generating key pairs in 193
ownership change privileges, file 52

P

panics 284
password rules, changing 127
passwords
changing (passwd) 125
managing security using 123
passwords, changing (passwd) 126
perfmon, using to monitor performance 277
performance, monitoring with perfmon 277
permissions of default directories (/etc, /home) 58
plexes, displaying information about 251
private and public keys. *See* RSA key pair 193
privilege levels for Data ONTAP commands 22
privileges, file ownership change 52
protocols. *See* SSH (Secure Shell) protocol, SSL (Secure Sockets Layer) protocol
public-key encryption 183, 184
PuTTY for SecureAdmin 183

Q

qtrees, definition of 295
quota file, character coding for 59

R

RAID
displaying statistics (aggr status) 251
displaying statistics (sysconfig -r) 246
displaying statistics (vol status) 253
RAID, definition of 295
rebooting the system
from the console 91
remotely 91
using RLM 91
regeneration of certificates 201
reinitialization
of SSH 190
of SSL 201
Remote LAN Module
definition of 6
remote management 205
Remote Management Controller 5
Remote Shell access to storage system (rsh) 37
requirements
for Web browser 40
to manage storage system on NFS clients 46
RLM
admin mode command syntax 221
admin mode commands 219
advanced command syntax 223
advanced mode commands 223
AutoSupport messages 228
booting N5200 and N5500 systems from 91
command line interface (CLI) 218
description of 204
displaying information in admin mode 224
displaying information in advanced mode 223
features 205
firware update problems, troubleshooting 242
how to configure 208
logging in to 215
managing with Data ONTAP commands 208, 213
system console redirection feature 219
System Event Log 226
troubleshooting communication problems 238
troubleshooting configuration problems 238
troubleshooting connection problems 239
troubleshooting firmware update problems 242
using AutoSupport options 214
RMC 5

roles
assigning capabilities to 109
assigning to groups 106
creating 111
definition of 98
deleting 119
listing 114
modifying 112
naming requirements 98
predefined 108
root option for aggregates 66
root password, changing 125
root volume
changing 65
default name 55
directories contained within 58
minimum size 56
size requirement 56
space guarantees and 57
RSA key pair
definition of 193
generating for SSH 1.x 193
generating for SSH 2.0 194
where stored 193, 194
RSA/DSA, for SecureAdmin 183
rsh (Remote Shell)
access to storage system 37
using with Windows 39
rsh commands
accessing storage system from a PC client 39
accessing storage system from a UNIX client 38
format used with user name and password 38
list of 38
privilege levels 22
use with user names and passwords 37

S

secure connection, testing 200
Secure FilerView, improving security using 186
secure session, creating with SecureAdmin 184
SecureAdmin
authentication supported 183
creating a secure session with 184
displaying status of 202
encryption algorithims supported 183
improving security with SSH 183
improving security with SSL 186
managing SSH portion 187
managing SSL portion 197
security
controlling file ownership changes (options wafl.root_only_chown) 52
controlling mount privileges (options nfs.mount_rootonly) 51
improving using Secure FilerView 186
improving using SecureAdmin 182
improving using SSH 183
limiting Remote Shell access 50
limiting Telnet access 49
options 53
password options 127
passwords, managing 123
self-signed certificates 197
serial ports (storage system hardware) 6
server keys
changing the size of 190
setting 188
size guidelines for 188
uses of 183
using with SSH 188
server requirements for netboot 88
session keys, uses of 184
share, definition of 295
single partitions 68
SnapLock 11
SnapMirror 11
SnapRestore 12
Snapshot copy software 12
Snapshot copy, aggregate 132
Snapshot copy, definition of 296
SnapVault 12
SSH (Secure Shell) commands
secureadmin disable all 202
secureadmin disable ssh 190, 192
secureadmin enable all 202
secureadmin enable ssh 188, 190, 192
secureadmin setup -f ssh 190
secureadmin setup ssh 188

secureadmin status 202
SSH (Secure Shell) protocol
authentication with 183
creating a secure session with 184
determining host and server key size using 188
disabling or enabling 188, 192
encryption with 183
host keys. *See* host keys
improving security with 183
keys used by 183
managing 187
reinitializing 190
server keys. *See* server keys
session keys 184
setting up and starting 188
SSH Communications Security client, for SecureAdmin 183
SSL (Secure Sockets Layer) commands
secureadmin addcert ssl 197, 201
secureadmin disable all 202
secureadmin disable ssl 201
secureadmin enable all 202
secureadmin enable ssl 197
secureadmin enable ssl (enables SSL) 201
secureadmin setup -f ssl 201
secureadmin setup ssl 197
secureadmin status 202
SSL (Secure Sockets Layer) protocol
authentication with 186
certificates used with 197
disabling or enabling 201
improving security with 186
managing 197
reinitializing 201
setting up and starting 197
startup commands in /etc/rc 152
statistics commands
aggr status command, description of 251
checking expansion cards 246
displaying adapter information 249
displaying aggregate state statistics 251
displaying chassis environment status 263
displaying Data ONTAP version 246
displaying disk information
aggr status 251
storage show 249
sysconfig -d 246
vol status 253
displaying Fibre Channel driver statistics 265
displaying link statistics 265
displaying medium changer information 249
displaying overall storage system information 247
displaying RAID and checksum information 246, 251, 253
displaying relative environment information 262
displaying relative physical drive position 265
displaying shelf environment status 263
displaying tape drive information 247
displaying tape information 250
displaying tape library information 247
displaying volume
information (sysconfig -v) 247
language (vol status) 254
displaying volume state statistics 253
environment, description of 262
Fibre Channel statistics, description of 265
storage command, description of 249
sysconfig command, description of 246
vol status command, description of 253
stats command
about 266
background mode 273
controlling output 274
counters 266
instances 266
objects 266
preset files 275
repeat mode 272
singleton mode 271
status commands
aggr status (displays aggregate state) 251
aggr status -d (displays disk statistics) 251
aggr status -r (displays RAID statistics) 251
environment chassis (displays shelf environment information) 263
environment command, description of 262
environment status (displays all storage system environment information) 262

environment status shelf (displays shelf environment information) 263
fcstat device_map (displays relative physical drive position) 265
fcstat fcal_stats (displays fibre channel driver statistics) 265
fcstat link_stats (displays link statistics) 265
filestats command, description of 255
storage command, description of 249
storage hub (displays hub information) 249
storage show adapter (displays adapter information) 249
storage show disk (displays disk information) 249
storage show mc (displays medium changer information) 249
storage show port (displays switch port information) 249
storage show switch (displays switch information) 249
storage show tape (displays tape information) 250
storage stats tape (displays tape statistics) 250
sysconfig -c (checks expansion cards) 246
sysconfig -d (displays disk information) 246
sysconfig -m (displays tape drive information) 247
sysconfig -m (displays tape library statistics) 247
sysconfig -r (displays RAID information) 246
sysconfig -v (displays overall filr statistics) 247
sysconfig -V (displays volume statistics) 247
ups (displays UPS environment information) 263
vol status (displays volume state) 253
vol status -d (displays disk statistics) 253
vol status -l (displays volume language) 254
vol status -r (displays RAID statistics) 253
status, displaying SecureAdmin 202
storage system
components of 3
definition of 2
management capabilities 13
storage system access 45
/etc directory, accessing from CIFS client 61
/etc directory, accessing from NFS client 61
/home directory, accessing from CIFS client 62
/home directory, accessing from NFS client 62
from the console 29, 30
inaccessiblility, reasons for 155
using FilerView 40
using rsh command from a PC client 39
using rsh command from a UNIX client 38
with Remote Shell (rsh) 37
with Telnet 33, 34
storage system hardware
board 4
components 4
environmental adapter 5
expansion slots 6
LCD 5
memory 4
serial ports 6
storage systems
booting 75
booting the system (boot) 68
changing domain name of 200
decryption 184
disabling licenses for 137
displaying current licenses for (license) 137
displaying overall statistics for 247
displaying version of 246
editing boot configuration file in 154
enabling licenses for (license) 136
halting (halt) 94
improving performance in 278
managing CIFS using FilerView 17
performance improvements
aggregate sizing 279
avoiding inodes access time update 278
backup rate 281
balancing NFS traffic on interfaces 278
caching client data to reduce traffic 280
disk-bound volume 279
large transfer 280
maintain adequate free blocks and inodes 279
read-ahead 279
reallocate command 280

using TCP 278
rebooting the system (reboot) 91
rsh (Remote Shell) access to 37
savecore (saves the core file) 143
security
controlling file ownership changes (options wafl.root_only_chown) 52
controlling mount privileges (options nfs.mount_rootonly) 51
limiting Telnet access 49
options 53
setting date and time (date) 138
support for AutoSupport, mail host 165
SyncMirror 13
syncronizing system time 140
Syslog Translator
accessing through FilerView 286
using 285
system
date and time, setting 138
panics 284
password, changing 125
rebooting, from the console 91
rebooting, remotely 91

T

tape backup and restore 13
tape drives, displaying statistics 247
tapes, displaying tape library statistics 247
TCP, increasing window size 280
Telnet
access to storage system 33
configuring a timeout period 35
limiting access to 49
termination of session with the storage system 34
TFTP
enabling console logging of accessed files 90
root directory, specifying 89
services, configuring 89
time
setting storage system date and 138
synchronizing 140
time servers, about 140
timed, using 140
troubleshooting
reasons to boot from diskette 85
troubleshooting, reasons to boot from diskette 70

U

Uninterruptible Power Supply (UPS)
adding a device to be monitored 263
enabling or disabling monitoring of 264
UPS
management 159
shutdown process 159
use of administration hosts 45
use of FilerView interface 42
user account, changing password for 126
useradmin
examples 120
naming requirements 98
usermap file, character coding for 59
users
assigning to groups 100
capabilities and 98
changing passwords 126
creation examples 120
definition of 97
deleting 119
examples of creating 120
listing 114
modifying capabilities of 104
naming requirement 98

V

Vandyke SecureCRT, for SecureAdmin 183
version checking, Data ONTAP 81
vFiler unit 14
virus scan support 13
vol copy 13
volumes
disk statistics, displaying (vol status) 253
vol status command, description of 253
volume language, displaying (vol status) 254
volume state, displaying (vol status) 253
volume statistics, displaying 247

W

WAFL 2
warnings
    man-in-the-middle 191
    obsolete domain names 200
warningtime (UPS option) 159
Windows
    administrator accounts in 97
    domain users, granting access to 102
    network commands 17
    special groups 99
    using rsh with 39