

Data ONTAP® 8.0 7-Mode

System Administration Guide

NetApp, 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: 210-05000_A0
Updated for Data ONTAP 8.0.1 on 04 November 2010

Contents

Copyright information	13
Trademark information	15
About this guide	17
Audience	17
Accessing Data ONTAP man pages	18
Terminology	18
Where to enter commands	19
Keyboard and formatting conventions	19
Special messages	20
How to send your comments	21
Introduction to NetApp storage	23
Components of a storage system	23
Internal components	24
Slots and ports	25
Disk shelves and disks	26
Third-party storage	26
Data ONTAP features	27
Network file service	27
Multiprotocol file and block sharing	28
Data storage management	28
Data organization management	28
Data access management	29
Data migration management	29
Data protection	29
System management	32
AutoSupport	33
How to interface with Data ONTAP	35
Methods for administering a storage system	35
Data ONTAP command-line interface	37
Using the history feature	37
Using the command-line editor	37
How to use online command-line help	38

Data ONTAP commands at different privilege levels	39
How different privilege settings apply to different sessions	39
Initial privilege level	39
Setting the privilege level	40
How to access the storage system	41
Methods for accessing a storage system	41
Methods for administering the system (no licenses are required)	41
Methods for storing and retrieving data (licenses are required)	42
Controlling the sharing of a console session	42
Rules that apply to console, Telnet, and SSH-interactive sessions	43
The eOM interface	45
How to access a storage system from the console	47
Using the serial port to access the storage system	47
Using the remote management device to access the system console	48
Secure protocols and storage system access	49
The default security settings	50
The SSH protocol	51
The SSL protocol	61
Determining whether secure protocols are enabled	66
Enabling or disabling secure protocols	66
How to access a storage system by using Telnet	66
Starting a Telnet session	67
Terminating a Telnet session	68
Configuration for Telnet sessions	68
How to access a storage system by using a Remote Shell connection	70
When to use RSH commands with user names and passwords	71
Accessing a storage system from a UNIX client by using RSH	72
Accessing a storage system from a Windows client by using a Remote Shell application	73
Commands not accepted when using RSH	73
How to reset options to default values from RSH	74
Displaying RSH session information	74
How to access a storage system by using FilerView	75
Accessing a storage system from a client by using FilerView	76
The FilerView interface	78
Read-only access to FilerView	79

How to manage access from administration hosts	79
Reasons to designate a workstation as an administrative host	79
Administration host privileges	80
Requirements for using a client	80
How to specify administration hosts	80
Adding administration hosts	81
Removing administration hosts	81
Methods for controlling storage system access	82
Controlling Telnet access using host names	82
Restricting protocol access	83
Controlling mount privilege	84
Controlling file ownership change privileges	84
Controlling anonymous CIFS share lookups	85
Options that help maintain security	85
Allowing only secure access to the storage system	87
How to manage the root volume	89
Recommendations regarding the root volume	89
Size requirement for root FlexVol volumes	91
Default directories in the root volume	92
Permissions for the default directories	92
The /etc directory	93
How to access the default directories on the storage system	97
Accessing the /etc directory from an NFS client	97
Accessing the /etc directory from a CIFS client	97
Accessing the /etc directory with FTP	98
Accessing the /home directory from an NFS client	98
Accessing the /home directory from a CIFS client	99
Accessing the /home directory with FTP	99
Accessing log files using HTTP or HTTPS	100
Changing the root volume	100
How to start and stop the storage system	103
How to boot the storage system	103
Ways to boot the storage system	103
Booting the storage system at the storage system prompt	104
Booting Data ONTAP at the boot environment prompt	107
Booting Data ONTAP remotely	108

Recovering from a corrupted image of the boot device	109
Checking available Data ONTAP versions	110
About rebooting the storage system	110
Rebooting the storage system from the system console	111
Rebooting the storage system remotely	111
Halting the storage system	112
How to manage administrator and diagnostic access	115
Reasons for creating administrator accounts	115
What users, groups, roles, and capabilities are	116
How users are assigned capabilities	117
Requirements for naming users, groups, and roles	117
Windows special groups	117
About changing capabilities of other groups and roles	118
Root access to the storage system	118
Disabling root access to the storage system	119
Displaying the status of root access	119
How to manage users	120
Creating users and assigning them to groups	120
Granting access to Windows domainusers	122
How to grant permissions for MMC	123
About changing another user's capabilities	123
How to manage groups	123
Predefined groups	124
Assigning roles to groups by creating or modifying a group	125
Renaming a group	126
Loading groups from the lclgroups.cfg file	126
Setting the maximum number of auxiliary UNIX groups allowed for a user	127
How to manage roles	127
Predefined roles	128
Supported capability types	129
Creating a new role and assigning capabilities to roles	131
Modifying an existing role or its capabilities	132
Users, groups, and roles	133
Commands that list users, domainusers, groups, or roles	133
Commands that delete users, domainusers, groups, or roles	137

Administrative user creation examples	138
Example of creating a user with custom capabilities	138
Example of creating a user with no administrative capabilities	139
How to manage passwords for security	140
Changing the storage system password	141
Changing a local user account password	142
Options that manage password rules	143
The diagnostic account and the systemshell	146
Enabling and disabling the diagnostic account	146
Setting the password for the diagnostic account	147
Accessing the systemshell	148
General system maintenance	151
Special system files	151
Aggregate Snapshot copy management	151
How to create aggregate Snapshot copies	152
Aggregate Snapshot reserve	152
Automatic aggregate Snapshot copy deletion	153
Disabling automatic aggregate Snapshot copy creation	154
Ways to manage licenses	154
Adding a license	155
Displaying current license codes	155
Disabling a license	155
Setting the system date and time	156
Synchronizing the system time	157
The timed options	157
Displaying and setting the system time zone	158
Core files	159
Core dump writing	159
Automatic technical support notification upon system reboots	160
Message logging	160
The /etc/syslog.conf file	161
Sample /etc/syslog.conf file	162
Configuring message logging	162
Audit logging	163
Configuring audit logging	164
Enabling or disabling read-only API auditing	164

Startup configuration for the storage system	165
About the /etc/rc file	165
Editing the /etc/rc file	166
Recovering from /etc/rc errors	167
Storage system configuration backup and cloning	168
Backing up a storage system configuration	168
Cloning a storage system configuration	169
Restoring a storage system configuration	169
Comparing storage system configurations and backup configuration files	170
About writing and reading files on the storage system	170
Writing a WAFL file	171
Reading a WAFL file	172
UPS management	172
The UPS shutdown options	173
The UPS shutdown process	173
Factors that might influence UPS shutdown event timing	174
The AutoSupport feature	175
Overview of the AutoSupport feature	175
AutoSupport transport protocols	176
Configuring AutoSupport	177
AutoSupport options	177
Testing AutoSupport	182
AutoSupport troubleshooting tasks	183
Troubleshooting AutoSupport over HTTP or HTTPS	183
Troubleshooting AutoSupport over SMTP	183
Controlling the size of AutoSupport messages	184
AutoSupport messages	185
Getting AutoSupport message descriptions	185
Contents of AutoSupport event messages	186
Managing storage systems remotely	189
Using the Service Processor for remote system management	189
Ways to configure the SP	191
Prerequisites for configuring the SP	191
Configuring the SP	192
Accounts that can access the SP	194

Logging in to the SP from an administration host	195
Accessing the SP from the system console	196
SP CLI and system console sessions	197
How to use the SP CLI	197
How to use Data ONTAP to manage the SP	203
How the SP sensors help you monitor system components	206
SP commands for troubleshooting the storage system	211
System event log and the SP	212
Console log and the SP	213
AutoSupport messages for systems with the SP	213
How to update the SP firmware	214
Troubleshooting SP connection problems	215
Using the Remote LAN Module for remote system management	216
What the RLM does	217
Ways to configure the RLM	218
How to log in to the RLM	221
How to manage the storage system with the RLM	225
How to manage the RLM with Data ONTAP	231
How to display information about the storage system and the RLM	234
Comparison of Data ONTAP and RLM commands	239
How to troubleshoot the storage system with the RLM	241
How to update the RLM firmware	241
How to troubleshoot RLM problems	242
Using the Baseboard Management Controller for remote system management	245
What the BMC does	247
Ways to configure the BMC	248
How to manage the BMC with Data ONTAP	251
How to log in to the BMC	253
How to manage the storage system with the BMC	255
How to display information about the storage system and the BMC	261
Comparison of Data ONTAP and BMC commands	266
How to troubleshoot the storage system with the BMC	268
How to update the BMC firmware	268
How to troubleshoot BMC problems	269
The Remote Support Agent as a firmware upgrade	273
System information	275

Getting storage system configuration information	275
Commands to display storage subsystem information	277
Getting aggregate information	279
Getting volume information	280
Getting a file statistics summary	281
Example of the filestats command with no options specified	283
Examples of the filestats command with ages option specified	284
Example of the filestats command with sizes option specified	285
Example of using the filestats command to determine volume capacity	285
Storage system environment information	285
Getting environmental status information	286
Specifying a UPS device to be monitored	287
Enabling or disabling monitoring of UPS devices	287
Getting Fibre Channel information	288
Getting SAS adapter and expander information	288
Storage system information and the stats command	289
Viewing the list of available counters	290
Getting detailed information about a counter	291
Using the stats command interactively in singleton mode	292
Using the stats command interactively in repeat mode	293
Collecting system information by using the stats command in background mode	294
Changing the output of a stats command	295
About the stats preset files	297
How to get system information using perfmon	298
How to get system information using perfstat	298
System performance and resources	299
How to manage storage system resources by using FlexShare	299
When to use FlexShare	299
How to use FlexShare	302
How to increase WAFL cache memory	307
Enabling and disabling WAFL extended cache	308
Caching normal user data blocks	308
Caching low-priority user data blocks	308
Caching only system metadata	309

Integrating FlexShare buffer cache policies with WAFL extended cache options	310
Displaying the WAFL extended cache configuration	311
Displaying usage and access information for WAFL extended cache	311
Ways to improve storage system performance	312
About balancing NFS traffic on network interfaces	312
How to ensure reliable NFS traffic by using TCP	312
Avoiding access time update for inodes	313
Adding disks to a disk-bound aggregate	313
About sizing aggregates appropriately	314
About putting cards into the correct slots	314
Maintaining adequate free blocks and free inodes	314
About optimizing LUN, file, and volume layout	315
Using oplocks for CIFS storage systems	315
Increasing the TCP window size for CIFS or NFS	315
About backing up by using qtrees	316
How to optimize LUN, file, volume, and aggregate layout	316
What a reallocation scan is	317
Reasons to use LUN, file, or volume reallocation scans	318
Reasons to use aggregate reallocation scans	318
Reasons to use physical reallocation scans	318
How a reallocation scan works	319
How you manage reallocation scans	320
How to use reallocation scans most efficiently	329
How to improve read performance	330
About read reallocation	330
About improving Microsoft Exchange read performance	331
Troubleshooting tools	335
Storage system panics	335
Reacting to storage system panics	335
Error messages	336
Using the Syslog Translator to get more information about error messages	336
Accessing the Syslog Translator using FilerView	337
How to use the NOW site for help with errors	337
How to use the remote management device to troubleshoot the system	338

Glossary	339
Index	345

Copyright information

Copyright © 1994–2010 NetApp, 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.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP 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.

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

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

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(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 NetApp logo; the Network Appliance logo; Bycast; Cryptainer; Cryptoshred; DataFabric; Data ONTAP; Decru; Decru DataFort; FA Server; FilerView; FlexCache; FlexClone; FlexShare; FlexVol; FPolicy; gFiler; Go further, faster; Manage ONTAP; MultiStore; NearStore; NetCache; NOW (NetApp on the Web); ONTAPI; RAID-DP; SANscreen; SecureShare; Simulate ONTAP; SnapCopy; SnapDrive; SnapLock; SnapManager; SnapMirror; SnapMover; SnapRestore; SnapValidator; SnapVault; Spinnaker Networks; Spinnaker Networks logo; SpinAccess; SpinCluster; SpinFlex; SpinFS; SpinHA; SpinMove; SpinServer; SpinStor; StorageGRID; StoreVault; SyncMirror; Topio; vFiler; VFM; and WAFL are registered trademarks of NetApp, Inc. in the U.S.A. and/or other countries. Network Appliance, Snapshot, and The evolution of storage are trademarks of NetApp, Inc. in the U.S.A. and/or other countries and registered trademarks in some other countries. The StoreVault logo, ApplianceWatch, ApplianceWatch PRO, ASUP, AutoSupport, ComplianceClock, DataFort, Data Motion, FlexScale, FlexSuite, Lifetime Key Management, LockVault, NOW, MetroCluster, OpenKey, ReplicatorX, SecureAdmin, Shadow Tape, SnapDirector, SnapFilter, SnapMigrator, SnapSuite, Tech OnTap, Virtual File Manager, VPolicy, and Web Filer are trademarks of NetApp, Inc. in the U.S.A. and other countries. Get Successful and Select are service marks of NetApp, Inc. in the U.S.A.

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

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the U.S.A. and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the U.S.A. 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 U.S.A. and/or other countries.

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

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

About this guide

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

This document describes how to configure, operate, and manage storage systems that run Data ONTAP software.

Note: This guide applies to systems running Data ONTAP 8.x 7-Mode, including V-Series systems. The *7-Mode* in the *Data ONTAP 8.x 7-Mode* product name means that this release has the features and functionality you are used to if you have been using the Data ONTAP 7.0, 7.1, 7.2, or 7.3 release families. If you are a Data ONTAP 8.x Cluster-Mode user, you use the Data ONTAP 8.x Cluster-Mode guides plus any Data ONTAP 8.x 7-Mode guides for functionality you might want to access with 7-Mode commands through the nodeshell.

Next topics

[Audience](#) on page 17

[Accessing Data ONTAP man pages](#) on page 18

[Terminology](#) on page 18

[Where to enter commands](#) on page 19

[Keyboard and formatting conventions](#) on page 19

[Special messages](#) on page 20

[How to send your comments](#) on page 21

Audience

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

This document is for system administrators who are familiar with operating systems such as UNIX® and Windows® that run on the storage system's clients.

This document assumes that you are familiar with how to configure the storage system and how Network File System (NFS), Common Internet File System (CIFS), Hypertext Transport Protocol (HTTP), File Transport Protocol (FTP), and Web-based Distributed Authoring and Versioning (WebDAV) are used for file sharing or transfers. This guide doesn't cover basic system or network administration topics, such as IP addressing, routing, and network topology; it emphasizes the characteristics of the storage system.

Accessing Data ONTAP man pages

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

About this task

Data ONTAP manual pages are available for the following types of information. They are grouped into sections according to standard UNIX naming conventions.

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

Step

1. View man pages in the following ways:

- Enter the following command at the console command line:
`man command_or_file_name`
- Click the manual pages button on the main Data ONTAP navigational page in the FilerView user interface.

Note: All Data ONTAP 8.x 7-Mode man pages are stored on the system in files whose names are prefixed with the string "na_" to distinguish them from other man pages. The prefixed names sometimes appear in the NAME field of the man page, but the prefixes are not part of the command, file, or service.

Terminology

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

array LUN The storage that third-party storage arrays provide to storage systems running Data ONTAP software. One array LUN is the equivalent of one disk on a native disk shelf.

LUN (logical unit number) A logical unit of storage identified by a number.

native disk	A disk that is sold as local storage for storage systems that run Data ONTAP software.
native disk shelf	A disk shelf that is sold as local storage for storage systems that run Data ONTAP software.
storage controller	The component of a storage system that runs the Data ONTAP operating system and controls its disk subsystem. Storage controllers are also sometimes called <i>controllers</i> , <i>storage appliances</i> , <i>appliances</i> , <i>storage engines</i> , <i>heads</i> , <i>CPU modules</i> , or <i>controller modules</i> .
storage system	The hardware device running Data ONTAP that receives data from and sends data to native disk shelves, third-party storage, or both. Storage systems that run Data ONTAP are sometimes referred to as <i>filers</i> , <i>appliances</i> , <i>storage appliances</i> , <i>V-Series systems</i> , or <i>systems</i> .
third-party storage	The back-end storage arrays, such as IBM, Hitachi Data Systems, and HP, that provide storage for storage systems running Data ONTAP.

Where to enter commands

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

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

- You can enter commands on the system console, or from any client computer that can obtain access to the storage system using a Telnet or Secure Socket Shell (SSH) session. In examples that illustrate command execution, the command syntax and output might differ, depending on your version of the operating system.
- You can use the FilerView graphical user interface.

Keyboard and formatting conventions

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

Keyboard conventions

Convention	What it means
The NOW site	Refers to the NetApp Support site at now.netapp.com .

Convention	What it means
<i>Enter, enter</i>	<ul style="list-style-type: none"> Used to refer to the key that generates a carriage return; the key is named Return on some keyboards. Used to mean pressing one or more keys on the keyboard and then pressing the Enter key, or clicking in a field in a graphical interface and then typing information into the field.
hyphen (-)	Used to separate individual keys. For example, Ctrl-D means holding down the Ctrl key while pressing the D key.
type	Used to mean pressing one or more keys on the keyboard.

Formatting conventions

Convention	What it means
<i>Italic font</i>	<ul style="list-style-type: none"> Words or characters that require special attention. Placeholders for information that you must supply. For example, if the guide says to enter the <code>arp -d hostname</code> command, you enter the characters "arp -d" followed by the actual name of the host. Book titles in cross-references.
Monospaced font	<ul style="list-style-type: none"> Command names, option names, keywords, and daemon names. Information displayed on the system console or other computer monitors. Contents of files. File, path, and directory names.
Bold monospaced font	Words or characters you type. What you type is always shown in lowercase letters, unless your program is case-sensitive and uppercase letters are necessary for it to work properly.

Special messages

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

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

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

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by e-mail to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the name of your product and the applicable operating system. For example, *FAS6070—Data ONTAP 7.3*, or *Host Utilities—Solaris*, or *Operations Manager 3.8—Windows*.

Introduction to NetApp storage

NetApp storage systems are hardware- and software-based data storage and retrieval systems. They respond to network requests from clients and fulfill them by writing data to or retrieving data from the disk arrays. They 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 the [NetApp Products and Technologies](#) page.

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 management device such as the Service Processor (SP), the Remote LAN Module (RLM), or the Baseboard Management Controller (BMC)—through the remote management device's Ethernet connection to the system console.

Next topics

[Components of a storage system](#) on page 23

[Data ONTAP features](#) on page 27

Related information

[The NetApp Products and Technologies page - www.netapp.com/products/](http://www.netapp.com/products/)

Components of a storage system

A storage system running Data ONTAP has a main unit, which is the hardware device that receives and sends data. Depending on the platform, a storage system uses storage on disk shelves, third-party storage, or both.

The storage system running Data ONTAP consists of the following components:

- The storage system main unit, or chassis, is also known as the storage engine. It is 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 environmental error codes, see the *Diagnostics Guide* on the NOW site.
- The disk shelves are the 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.

Note: For V-Series systems, back-end storage arrays such as IBM, Hitachi Data Systems, and HP provide storage for data. V-Series systems fulfill client requests from either disk shelves or logical unit numbers (LUNs) on the back-end storage arrays.

More specifically, the storage system includes internal components, slots and ports, and disk shelves that contain the disks.

Next topics

Internal components on page 24

Slots and ports on page 25

Disk shelves and disks on page 26

Third-party storage on page 26

Related concepts

Storage system environment information on page 285

Related information

The NOW site - <http://now.netapp.com/>

Internal components

The internal components of a storage system enable the system to function.

The following table shows the internal components of a storage system.

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.
boot device	The storage system automatically boots from a Data ONTAP release stored on the boot device, such as a PC CompactFlash card. The boot device also stores a backup version of Data ONTAP from which to boot the storage system in an emergency.
LCD and LEDs	The storage system displays status information on the LCD and LEDs.

Component	Description
environmental adapter	The environmental adapter performs the following functions: <ul style="list-style-type: none"> • 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.
The remote management device such as the Service Processor (SP), the Remote LAN Module (RLM), or the Baseboard Management Controller (BMC)	The remote management device provides remote platform management capabilities for the storage system, allowing 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 remote management device monitors and maintains hardware event logs for the storage system and generates alerts based on system status.

Related concepts

[Using the Remote LAN Module for remote system management](#) on page 216

[Using the Baseboard Management Controller for remote system management](#) on page 245

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*.

The following table describes the slots and ports of a storage system.

Component	Description
slots	The storage system contains expansion slots for the following host adapters: <ul style="list-style-type: none"> • Network interface cards (NICs) • Adapters for the disk shelf or tape drive • Performance Acceleration Modules • NVRAM adapters

Component	Description
serial ports	<p>The serial ports include:</p> <ul style="list-style-type: none"> • The console port, which connects the storage system to a serial terminal that you can use as a console. • The port for remote management or diagnostics, which can be used for Data ONTAP management activities or connects diagnostic equipment, such as the environmental monitor unit (EMU) of a 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 *Data ONTAP 7-Mode Storage Management Guide*.

For information about disk shelves connected to V-Series systems, see the *V-Series Systems Implementation Guide for Native Disk Shelves* and the disk shelf guide.

Third-party storage

On a V-Series system, Data ONTAP provides unified NAS and SAN access to data stored in heterogeneous Fibre Channel (FC) SAN storage arrays, including storage arrays from IBM, Hitachi Data Systems, HP, and EMC. Data ONTAP supports multiple storage arrays of the same model or different models behind one V-Series system.

The Data ONTAP software provides a unified storage software platform that simplifies managing LUNs on storage arrays and storage on disk shelves. You can add storage when and where you need it, without disruption.

For information about supported storage array models, see the *V-Series Support Matrix*.

For information about setting up a specific storage array to work with Data ONTAP, see the V-Series *Implementation Guides*.

Data ONTAP features

Data ONTAP provides features for network file service, multiprotocol file and block sharing, data storage management, data organization management, data access management, data migration management, data protection system management, and AutoSupport.

Next topics

Network file service on page 27

Multiprotocol file and block sharing on page 28

Data storage management on page 28

Data organization management on page 28

Data access management on page 29

Data migration management on page 29

Data protection on page 29

System management on page 32

AutoSupport on page 33

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 *Data ONTAP 7-Mode Network Management Guide*.

For information about configuring a storage system in a storage area network (SAN) fabric, see the NetApp Interoperability Matrix and the *Data ONTAP 7-Mode Block Access Management Guide for iSCSI and FC*.

Related information

NetApp Interoperability Matrix - <http://now.netapp.com/NOW/products/interoperability/>

Multiprotocol file and block sharing

Several protocols allow you 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 *Data ONTAP 7-Mode File Access and Protocols Management Guide*.

For information about SAN block access protocols, see the *Data ONTAP 7-Mode Block Access Management Guide for iSCSI and FC*.

Data storage management

Data ONTAP stores data on disks in disk shelves connected to storage systems or uses storage on third-party storage arrays.

For native storage, Data ONTAP uses RAID-DP or RAID4 groups to provide parity protection. For third-party storage, Data ONTAP uses RAID0 groups to optimize performance and storage utilization. The storage arrays provide the parity protection for third-party storage. Data ONTAP RAID groups are organized into plexes, and plexes are organized into aggregates.

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 Logical Unit Numbers (LUNs) in SAN environments. Aggregates provide the physical storage to contain volumes.

For more information, see the *Data ONTAP 7-Mode Storage Management Guide* and the *Data ONTAP 7-Mode Block Access Management Guide for iSCSI and FC*.

When Data ONTAP is installed on a storage system at the factory, a root volume is configured as `/vol/vol0`, which contains system files in the `/etc` directory.

Related concepts

[How to manage the root volume](#) on page 89

Data access management

Data ONTAP enables you to manage access to data.

Data ONTAP performs the following operations for data access management:

- Checks file access permissions against file access requests.
- Checks write operations against file and disk usage quotas that you set.
For more information, see the *Data ONTAP 7-Mode 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.
For more information on Snapshot copies, see the *Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide*.

Data migration management

Data ONTAP enables you to manages data migration.

Data ONTAP offers the following features for data migration management:

- Snapshot copies
- Asynchronous mirroring
- Synchronous mirroring
- Backup to tape
- Aggregate copy
- Volume copy
- FlexClone
- `ndmpcopy`

Data protection

Storage systems provide a wide range of data protection features such as aggr copy, MetroCluster, NDMP, NVFAIL, SnapMirror, SnapRestore, Snapshot, SnapVault, SyncMirror, Tape backup and restore, Virus scan support, and vol copy.

These features are described in the following table.

Feature	Description
aggr copy	<p>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.</p> <p>For information about aggregates and <code>aggr copy</code>, see the <i>Data ONTAP 7-Mode Storage Management Guide</i>.</p>
MetroCluster	<p>MetroCluster enhances SyncMirror functionality for disaster recovery by providing continuous volume mirroring over 500-meter to 30-kilometer distances.</p> <p>For information about disaster protection using MetroCluster, see the <i>Data ONTAP 7-Mode High-Availability Configuration Guide</i>.</p>
NDMP (Network Data Management Protocol)	<p>NDMP support enables third-party applications that use NDMP to manage tape backup operations of system data. The <code>ndmpcopy</code> command carries out NDMP-compliant backups and restores. Security login restricts access to NDMP operations.</p> <p>For information about NDMP, see the <i>Data ONTAP 7-Mode Data Protection Tape Backup and Recovery Guide</i>.</p>
NVFAIL	<p>The <code>nvfail</code> option provides protection against data corruption by nonvolatile RAM (NVRAM) failures.</p> <p>For information about NVFAIL, see the <i>Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide</i>.</p>
SnapMirror software (license required)	<p>System-to-system Snapshot 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.</p> <p>For information about SnapMirror, see the <i>Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide</i>.</p>
SnapRestore software (license required)	<p>The SnapRestore feature performs fast restoration of backed-up data on request from Snapshot copies on an entire volume.</p> <p>For information about SnapRestore, see the <i>Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide</i>.</p>

Feature	Description
Snapshot software	<p>Manual or automatically scheduled multiple backups (or Snapshot copies) of data using a minimal amount of additional disk space at no performance cost.</p> <p>For information about how Data ONTAP organizes and manages data, see the <i>Data ONTAP 7-Mode Storage Management Guide</i>.</p> <p>For information about Snapshot copies, see the <i>Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide</i>.</p>
SnapVault software (license required)	<p>SnapVault combines Snapshot schedules and Qtree SnapMirror to provide disk-based data protection for NetApp storage systems. You can also install the Open Systems SnapVault agent on non-NetApp systems. This allows SnapVault to back up and restore data to those systems also.</p> <p>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 do not need to perform full dumps from primary storage, so you do not need to schedule a backup window.</p> <p>For information about SnapVault, see the <i>Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide</i>.</p>
SyncMirror (high-availability configuration required)	<p>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. Data ONTAP supports RAID4 and RAID-DP only for disk shelves.</p> <p>Similarly, SyncMirror can be used for mirroring of third-party storage. In the case of an unrecoverable error, Data ONTAP automatically switches access to the mirrored plex on the other storage array. Data ONTAP uses RAID0 for managing storage on array LUNs, but the storage arrays provide RAID protection for third-party storage.</p> <p>For information about supported RAID levels and plexes, see the <i>Data ONTAP 7-Mode Storage Management Guide</i>. For information about SyncMirror, see the <i>Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide</i>.</p>

Feature	Description
Tape backup and restore	<p>Tape backup <code>dump</code> and <code>restore</code> 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.</p> <p>For information about tape backup, see the <i>Data ONTAP 7-Mode Data Protection Tape Backup and Recovery Guide</i>.</p>
Virus scan support	<p>Data ONTAP provides support for third-party-scanning software for files accessed by CIFS clients.</p> <p>For information about virus protection for CIFS, see the <i>Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide</i>.</p>
vol copy	<p>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.</p> <p>For information about volumes and <code>vol copy</code>, see the <i>Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide</i>.</p>

System management

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

You can use Data ONTAP to perform the following system management tasks:

- Manage network connections
- Manage adapters
- Manage protocols
- Configure a pair of storage systems into high-availability configuration for failover
- Configure SharedStorage storage systems into a community
- Manage storage and quotas
- 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 *Data ONTAP 7-Mode MultiStore Management Guide*

For information about all Data ONTAP commands, see the *Data ONTAP 7-Mode Commands: Manual Page Reference, Volume 1* and the *Data ONTAP 7-Mode Commands: Manual Page Reference, Volume 2*.

AutoSupport

AutoSupport automatically sends AutoSupport Mail notifications about storage system problems to technical support and designated recipients.

Related concepts

[The AutoSupport feature](#) on page 175

How to interface with Data ONTAP

You interface with Data ONTAP to administer your storage system.

Next topics

[Methods for administering a storage system](#) on page 35

[Data ONTAP command-line interface](#) on page 37

[Data ONTAP commands at different privilege levels](#) on page 39

Methods for administering a storage system

You can use Data ONTAP, the remote management device (the SP, the RLM, or the BMC), Windows, configuration files, FilerView, System Manager, the DataFabric Manager software, or the Manage ONTAP Developer SDK software to administer a storage system.

- **Command execution through the storage system's CLI**

The storage system's CLI enables you to execute all Data ONTAP administrative commands, with the exception of some Windows server administrative commands.

The Data ONTAP command line allows you to enter a maximum of 2,046 characters, and it supports a maximum number of 255 arguments for a single command.

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 the remote management device in the storage system
 - 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 remote management device**

The redirection feature of the remote management device enables you to remotely execute all Data ONTAP administrative commands.
- **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 can 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.

Note: For information on how to set up CIFS so that you can use a Windows client to access files on the storage system, see the *Data ONTAP 7-Mode Software Setup Guide*.

- **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.
- **System Manager**
System Manager provides setup and management capabilities for SAN and NAS environments from a Microsoft Windows system. You can use System Manager to quickly and efficiently set up storage systems that are single or in a high-availability configuration. You can also use System Manager to configure all protocols, such as NFS, CIFS, FCP and iSCSI, supply provisions for file sharing and applications, and monitor and manage your storage system. For more information about System Manager, see the NOW site.
- **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 on the NOW site.
- **Manage ONTAP SDK software**
Manage ONTAP SDK contains resources necessary to develop third-party applications which monitor and manage storage systems. The Manage ONTAP SDK kit is available to all NOW users for free download. It 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, which simulates the NetApp storage system to a very low level, is also available as a separate distribution. For more information, see the Manage ONTAP SDK page.

Related concepts

Managing storage systems remotely on page 189

Default directories in the root volume on page 92

Related information

The NOW site - <http://now.netapp.com/>

Manage ONTAP SDK - <http://communities.netapp.com/docs/DOC-1110>

Data ONTAP command-line interface

Data ONTAP provides several features to assist you when you enter commands on the command line.

When using the Data ONTAP command line, be aware of the following general rules:

- If you are entering a command with an element that includes a space, you must quote that element. For example,

```
toaster> environment status chassis "Power Supply"
```
- Do not use a # character in the command string.
 A # character always means to comment out the rest of the line, so Data ONTAP will ignore any information following the #.

Next topics

[Using the history feature](#) on page 37

[Using the command-line editor](#) on page 37

[How to use online command-line help](#) on page 38

Using the history feature

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

Step

1. Do one of the following:

If you want to...	Then...
Scroll back through commands	Press the Up arrow key or press Ctrl-P.
Scroll forward through commands	Press the Down arrow key or 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.

About this task

You can use various key combinations to move the cursor within the same line and edit the command, as shown in the following table.

Step

1. Do one of the following:

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

How to use online command-line help

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.

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).

The following example shows the result of entering the `environment help` command at the storage system command line. The command output displays the syntax help for the `environment` commands.

```
toaster> environment help
Usage: environment status |
[status] [shelf [<adapter>]] |
```

```
[status] [shelf_log] |
[status] [shelf_stats] |
[status] [shelf_power_status] |
[status] [chassis [all | list-sensors | Fan | Power | Temp | Power Supply
| RTC Battery | NVRAM4-temperature-7 | NVRAM4-battery-7]]
```

Related concepts

[Data ONTAP commands at different privilege levels](#) on page 39

Data ONTAP commands at different privilege levels

Data ONTAP provides two sets of commands, depending on the privilege level you set. The administrative level enables you to access commands that are sufficient for managing your storage system. The advanced level provides commands for troubleshooting, in addition to all the commands available at the administrative level.

Attention: Commands accessible only at the advanced level should be used under the guidance of technical support. Using some advanced commands without consulting technical support might result in data loss.

Next topics

[How different privilege settings apply to different sessions](#) on page 39

[Initial privilege level](#) on page 39

[Setting the privilege level](#) on page 40

How different privilege settings apply to different sessions

Sessions opened through the console, Telnet, and secure shell applications share the same privilege setting. However, you can set a different privilege level for each RSH invocation.

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.

However, if 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 a script invokes 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 in the subsequent RSH sessions, because the privilege level for each subsequent session is reset to administrative.

Setting the privilege level

You set the privilege level to access commands at either the administrative or the advanced level.

Step

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
```

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 technical personnel.
```

```
sys1*>
```

How to access the storage system

You can access the storage system from the console or through a Telnet session, a Remote Shell connection, a secure shell client application, or the FilerView.

Next topics

Methods for accessing a storage system on page 41

How to access a storage system from the console on page 47

Secure protocols and storage system access on page 49

How to access a storage system by using Telnet on page 66

How to access a storage system by using a Remote Shell connection on page 70

How to access a storage system by using FilerView on page 75

How to manage access from administration hosts on page 79

Methods for controlling storage system access on page 82

Methods for accessing a storage system

To access the storage system, you only need network connectivity to the storage system and authentication privileges, and no licenses are required. To store and retrieve data on the storage system, you must have an NFS or a CIFS license installed.

Next topics

Methods for administering the system (no licenses are required) on page 41

Methods for storing and retrieving data (licenses are required) on page 42

Controlling the sharing of a console session on page 42

Rules that apply to console, Telnet, and SSH-interactive sessions on page 43

The e0M interface on page 45

Methods for administering the system (no licenses are required)

You can access a storage system to administer it by using a serial console or through a NIC installed in the storage system.

These are the methods you can use, and no licenses are required:

- 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

Methods for storing and retrieving data (licenses are required)

You can access a storage system to administer it and to store and retrieve data, by using a serial console or through a NIC installed in the storage system.

These are the methods you can use, and licenses are required:

- 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

Note:

If you use the `wrfile` command to redirect input into non-interactive SSH, the command will fail if:

- SSH is configured to automatically send EOF's.
- SSH is used with the option `-n`, which sends EOF at the beginning of the message.
- 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)
- From an NFS client or CIFS client to access a LUN in a SAN environment by using the Internet SCSI (iSCSI) protocol or the Fibre Channel (FC) protocol.

Controlling the sharing of a console session

A console session can be shared with a Telnet or an SSH-interactive session at the same time, or it can be a distinct user environment, separate from Telnet and SSH-interactive sessions.

About this task

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. To enhance security, you should

ensure that the option is set to `on` to keep the console session separate from a Telnet or an SSH-interactive session.

The console session is always shared with the remote management device, regardless of the `telnet.distinct.enable` option setting.

Step

1. To control the sharing of a console session, enter the following command:

```
options telnet.distinct.enable [on|off]
```

Setting the option to `on` enhances security by keeping the console session separate from a Telnet or an SSH-interactive session. On storage systems shipped with Data ONTAP 8.0 or later, the default for this option is `on`.

Setting the option to `off` causes the console session to share with a Telnet or an SSH-interactive session. You cannot set the option to `off` if a user is currently assigned to the Compliance Administrators group.

If the `telnet.distinct.enable` option setting is changed during a Telnet or an SSH-interactive session, the change does not go into effect until the next Telnet or SSH login.

If you change the option setting after upgrading to Data ONTAP 8.0 or later, the changes are preserved even if the system reverts back to the previous Data ONTAP version.

Note: You can initiate an SSH-interactive session by opening the session without entering a command. For example, you would enter the following command:

```
ssh storage_system -l root:""
```

If you enter the following command instead, you would initiate a non-interactive session:

```
ssh storage_system -l root:"" command
```

Related concepts

[Options that help maintain security](#) on page 85

[Predefined groups](#) on page 124

[Predefined roles](#) on page 128

[Supported capability types](#) on page 129

Related tasks

[Creating users and assigning them to groups](#) on page 120

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

You cannot open both a Telnet and an SSH-interactive session at the same time. However, you can configure for the console to share a session with a Telnet or an SSH-interactive session.

The following rules apply to console, Telnet, and SSH-interactive sessions.

- Sharing the console session

If the `telnet.distinct.enable` option is set to `off`, the console shares a session with a Telnet or an 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.
- Audit-log 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
```

- Audit-log entries identify all Telnet and SSH-interactive commands as “telnet shell.”
- If the `autologout.telnet.enable` option is set to `on`, the autologout program logs the user out of the Telnet or SSH-interactive 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 10 minutes, every time you press Enter, the timeout counter starts counting. If 10 minutes elapse before you press Enter 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 among console, Telnet, and SSH-interactive sessions.
- Audit-log entries identify all console, Telnet, and SSH-interactive commands as “console shell.”
- If the `autologout.telnet.enable` option is set to `on`, the autologout program logs the user out of the Telnet or SSH-interactive session after the number of minutes specified by the `autologout.telnet.timeout` option has elapsed.

The timeout counter starts after the command is executed.

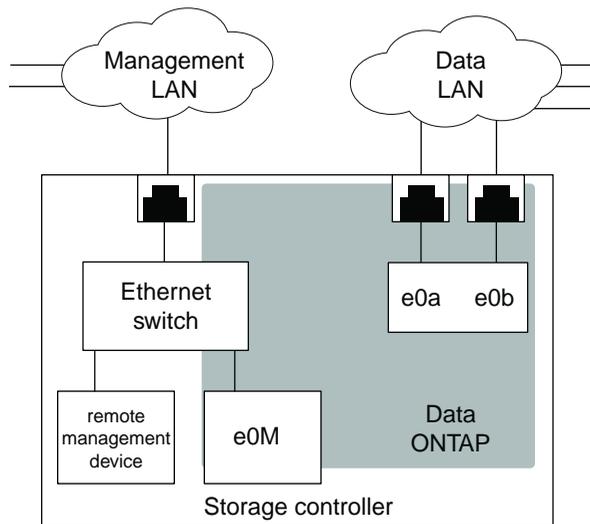
You can prevent commands from being aborted at the console or through a Telnet or an SSH session by using the `rsh` command to initiate commands from an administration host.

The `autologout.telnet.enable` and `autologout.telnet.timeout` options control the automatic timeout for both Telnet and SSH-interactive sessions. Even if you disable Telnet connections to the storage system, you can still enable and configure the automatic timeout period for only SSH-interactive sessions by setting the `autologout.telnet.enable` option to `on` and setting the `autologout.telnet.timeout` option to the desired timeout period.

The e0M interface

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

On a storage system that has the e0M interface, the Ethernet port (indicated by a wrench icon on the rear of the chassis) connects to an internal Ethernet switch. The internal Ethernet switch provides connectivity to the e0M interface and the remote management device such as the SP, the RLM, or the BMC. The following diagram illustrates the connections.



When you set up a system that includes the e0M interface, the Data ONTAP setup script recommends that you use the e0M as the preferred management interface for environments that use dedicated LANs to isolate management traffic from data traffic. The setup script then prompts you to configure e0M. The e0M configuration is separate from the configuration of the remote management device. Both configurations require unique IP addresses to allow the Ethernet switch to direct traffic to either the e0M interface or the remote management device. For information about how to set up the e0M interface, see the *Data ONTAP 7-Mode Software Setup Guide*.

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

- Telnet
- RSH
- HTTP or HTTPS
- SSH
- SNMP

Next topics

Using the e0M interface to perform a Data ONTAP management task on page 46

How the e0M interface and the remote management device differ on page 46

Related concepts

Using the Remote LAN Module for remote system management on page 216

Using the e0M interface to perform a Data ONTAP management task

You can use the e0M interface to access the storage system to manage Data ONTAP.

Steps

1. Open a Telnet, RSH, or SSH session on a client.

For information on how to use the e0M interface with SNMP, see the *Data ONTAP 7-Mode Network Management Guide*.

2. Connect to the storage system using the address of the e0M interface.
3. Log in to the storage system with an appropriate user name and a valid password.
4. At the storage system prompt, enter a Data ONTAP CLI command.

Example

To obtain the Data ONTAP version information, enter `version`.

Related concepts

How to access a storage system by using Telnet on page 66

How to access a storage system by using a Remote Shell connection on page 70

How to manage SSH on page 53

The default security settings on page 50

How the e0M interface and the remote management device differ

The e0M interface and the remote management device (which can be the SP, the RLM, or the BMC, depending on the storage system model) serve different functionality. Whereas the e0M interface serves as the dedicated interface for management traffic, the remote management device provides remote management capabilities.

The e0M interface serves as the dedicated interface for environments that have dedicated LANs for management traffic. You use the e0M interface for Data ONTAP administrative tasks.

The remote management device, on the other hand, not only can be used for managing Data ONTAP but also provides remote management capabilities for the storage system, including remote access to the console, monitoring, troubleshooting, logging, and alerting features. Also, the remote management device stays operational regardless of the operating state of the storage system and regardless of whether Data ONTAP is running or not.

Both the e0M interface and the remote management device connect to the internal Ethernet switch that connects to the Ethernet port. (The Ethernet port is indicated by a wrench icon on the rear of the chassis.)

How to access a storage system from the console

You can access the console to manage the storage system by using the serial port or a remote management device such as the SP, the RLM, or the BMC.

If you change the values of the following `options` commands, you must reestablish the console session before the values can take effect.

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

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

Next topics

[Using the serial port to access the storage system](#) on page 47

[Using the remote management device to access the system console](#) on page 48

Related concepts

[Rules that apply to console, Telnet, and SSH-interactive sessions](#) on page 43

Using the serial port to access the storage system

You can access a storage system directly from a console that is attached by a cable to the system's serial port.

Steps

1. At the console, press Enter.

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.

The storage system responds with the password prompt.

3. Enter the password for the root or administrative user account, or, if no password is defined, press Enter.
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 remote management device to access the system console

You can access a system console remotely by using the system console redirection feature provided by the remote management device. Depending on your storage system, the remote management device can be the SP, the RLM, or the BMC.

About this task

To log in to the SP or the RLM, you can use the naroot account or a Data ONTAP user account with the credentials of the admin role or a role with the `login-sp` capability.

To log into the BMC, you can use the root, naroot, or Administrator account.

Steps

1. From the administration host, log in to the remote management device by entering the following command:

```
ssh username@IP_for_remote_management_device
```

The storage system responds by displaying the CLI prompt for the remote management device.

2. Enter the following command at the CLI prompt for the remote management device:

```
system console
```

3. If the storage system displays the login prompt, enter an appropriate account name:

If you are using...	Enter the following account name...
The system root account	root

If you are using...	Enter the following account name...
An administrative user account	<i>username</i>
	Note: <i>username</i> is the administrative user account.

4. Enter the password for the account, or, if no password is defined, press Enter.
The storage system prompt appears.
5. To exit the console, do one of the following.
 - To exit the console redirection session and return to the SP prompt or the RLM prompt, press Ctrl-D.
 - To exit the console redirection session and return to the BMC prompt, press Ctrl-G.

Related concepts

[Managing storage systems remotely](#) on page 189

[Prerequisites for logging in to the SP](#) on page 0

[How to log in to the RLM](#) on page 221

[How to log in to the BMC](#) on page 253

Secure protocols and storage system access

Using secure protocols improves the security of your storage system by making 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.

If your storage system does not have secure protocols enabled, you can set up SecureAdmin, which provides a secure communication channel between a client and the storage system by using one or both of the following protocols—SSH and SSL.

Note: SecureAdmin is set up automatically on storage systems shipped with Data ONTAP 8.0 or later.

- Secure Shell (SSH) protocol
SSH provides a secure remote shell and interactive network session.
- Secure Sockets Layer (SSL) protocol
SSL provides secure Web access for FilerView and Data ONTAP APIs.

Next topics

[The default security settings](#) on page 50

[The SSH protocol](#) on page 51

[The SSL protocol](#) on page 61

[Determining whether secure protocols are enabled](#) on page 66

[Enabling or disabling secure protocols](#) on page 66

The default security settings

On storage systems shipped with Data ONTAP 8.0 or later, secure protocols are enabled and nonsecure protocols are disabled by default.

SecureAdmin is set up automatically on storage systems shipped with Data ONTAP 8.0 or later. For these systems, the following are the default security settings:

- Secure protocols (including SSH, SSL, and HTTPS) are enabled by default.
- Nonsecure protocols (including RSH, Telnet, FTP, and HTTP) are disabled by default.

On storage systems shipped with Data ONTAP 8.0 or later, the following are the default option settings for SSH and SSL:

- `options ssh.enable on`
- `options ssh2.enable on`
- `options ssh1.enable off`
- `options ssh.passwd_auth.enable on`
- `options ssh.pubkey_auth.enable on`
- `options httpd.admin.ssl.enable on`

Also on storage systems shipped with Data ONTAP 8.0 or later, the following are the default option settings for the nonsecure protocols:

- `options ftpd.enable off`
- `options httpd.admin.enable off`
- `options httpd.enable off`
- `options rsh.enable off`
- `options telnet.distinct.enable on`
- `options telnet.enable off`

Note: These default settings apply only to storage systems shipped with Data ONTAP 8.0 or later. For storage systems upgraded from an earlier version to Data ONTAP 8.0 or later, the above default settings do not apply. Instead, for those upgraded systems, the settings remain unchanged after the upgrade. Also, if you make security setting modifications after upgrading to Data ONTAP 8.0 or later, the modifications are preserved even if the system reverts back to the previous Data ONTAP version.

Related tasks

[Allowing only secure access to the storage system](#) on page 87

The SSH protocol

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.

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

Data ONTAP supports the following encryption algorithms:

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

Data ONTAP supports the SSH 1.x protocol and the SSH 2.0 protocol.

Data ONTAP supports the following SSH clients:

- OpenSSH client versions 3.8p1 and 4.4p1 on UNIX platforms
- SSH Communications Security client (SSH Tectia client) version 6.0.0 on Windows platforms
- Vandyke SecureCRT version 6.0.1 on Windows platforms
- PuTTY version 0.6.0 on Windows platforms
- F-Secure SSH client version 7.0.0 on UNIX platforms

SSH uses three keys to improve security:

- 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.

Note: SecureAdmin is set up automatically on storage systems shipped with Data ONTAP 8.0 or later.

- 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.

The following table shows how Data ONTAP 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.	
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.	

If you are logged into a non-root user account on a client, and you request a list of supported SSH commands on a storage system using the `ssh <ip address> ?` command, some SSH clients do not pass the ? (question mark) to the storage system. To make sure the client passes the question mark, wrap the ? in quotes, for example, `ssh <ip address> '?'`.

Note: Some characters, for example ?, ., *, and ^, can have special meaning for the command interpreter running on the client. The client command interpreter might replace the character with an environment-specific value prior to passing it to the SSH program. To prevent a replacement, use an escape sequence before the character (for example, `ssh <ip address> \?`) or enclose the character in quotes.

Next topics

[How to manage SSH](#) on page 53

[Setting up and starting SSH](#) on page 53

[Reinitializing SSH](#) on page 55

[Enabling or disabling SSH](#) on page 56

[Public-key-based authentication](#) on page 56

[Issuing SSH requests](#) on page 59

[Displaying the current SSH settings](#) on page 60

How to manage SSH

If your storage system does not have SSH enabled, you can set up SecureAdmin to enable secure sessions using SSH. A few options enable you to control password-based authentication and public key authentication, control access to a storage system, and assign the port number to a storage system.

SecureAdmin is set up automatically on storage systems shipped with Data ONTAP 8.0 or later.

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

- `options ssh.passwd_auth.enable`—Controls password-based authentication. The default is on.
- `options ssh.pubkey_auth.enable`—Controls public key authentication. The default is on.
- `options ssh.access`—Controls access to a storage system. The default value allows everyone to access the storage system.
- `options ssh.port`—Assigns the port number to a storage system. The default value is 22.

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

Note: SSH does not support `force` commands. It does not support internal role-based access control. Access control is governed by the Administrative Roles feature.

Related concepts

[How to manage administrator and diagnostic access](#) on page 115

[The default security settings](#) on page 50

Related tasks

[Restricting protocol access](#) on page 83

Setting up and starting SSH

The SSH 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.

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

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

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.

Steps

1. Enter the following command:

```
secureadmin setup [-f] [-q] ssh
```

The `-f` option forces setup to run even if the SSH server has already been configured.

The `-q` option is the non-interactive mode for setting up SSH. See the `na_secureadmin(1)` man page for more information.

2. When prompted, enter a size for the host key if you are using the SSH 1.x protocol.

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.

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.

The default size for the host key is 768 bits.

5. When prompted, confirm the parameters that you specified.

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

Reinitializing SSH enables you to change the sizes of existing host and server keys.

Steps

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 [-q] ssh
```

The `-f` option forces setup to run even if the SSH server has already been configured.

The `-q` option is the non-interactive mode for setting up SSH. See the `na_secureadmin(1)` man page for more information.

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.

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)?
```

Enabling or disabling SSH

After setting up SSH, you can enable or disable it to start or stop SSH service.

Step

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

```
secureadmin {enable|disable} {ssh1|ssh2}
```

Use `enable` to start SSH service or `disable` to stop SSH service.

Use `ssh1` to administer SSH 1.x clients or `ssh2` to administer SSH 2.0 clients.

Example of enabling SSH service for SSH 2.0 clients

The following command enables SSH service for SSH 2.0 clients:

```
secureadmin enable ssh2
```

Related tasks

[Setting up and starting SSH](#) on page 53

Public-key-based authentication

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.

Next topics

[Generating an RSA key pair for SSH 1.x](#) on page 56

[Generating key pairs for SSH 2.0](#) on page 57

[Editing public keys generated by SecureCRT and ssh.com clients](#) on page 58

Generating an RSA key pair for SSH 1.x

Public-key-based authentication using SSH 1.x requires an RSA key pair.

Steps

1. Using your SSH 1.x client, generate an RSA key pair.

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 of generating an RSA key pair

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.

Steps

1. Using your SSH 2.0 client, generate an RSA key pair.

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.

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_keys2` file.

Example of generating RSA and DSA key pairs

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_keys2` file on storage system `sys1`:

```
% mount sys1:/ /mnt_sys1
% cat id_rsa.pub >> /mnt_sys1/etc/sshd/john/.ssh/
authorized_keys2
% cat id_dsa.pub >> /mnt_sys1/etc/sshd/john/.ssh/
authorized_keys2
```

Related tasks

[Editing public keys generated by SecureCRT and ssh.com clients](#) on page 58

Editing 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 edit the generated public keys before SecureAdmin can use them.

Steps

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

Example of editing keys generated by SecureCRT

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"
AAAAB3NzaC1yc2EAAAADAQABAAQgQDJhJ6nk
+2hm5iZnx737ZqxfgksPl3+OY1cP80s
1amXuUrwBp3/MUODEP5E51lZqj00w5kyJlvPjCiLg9UqS7JeY5yd/
6xyGarsde26De1E
rbVJluqnxxyA0lV9A1hjBE8TbI+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 AAAAB3NzaC1yc2EAAAADAQABAAQgQDJhJ6nk
+2hm5iZnx737ZqxfgksPl
3+OY1cP80s1amXuUrwBp3/MUODEP5E51lZqj00w5kyJlvPjCiLg9UqS7JeY5yd/
6xy
Garsde26De1ErbVJluqnxxyA0lV9A1hjBE8TbI+lyYBH
+WezT0nySix6VBQTAWhv43r
9lSudswYV80Q==
```

Issuing SSH requests

You can issue SSH requests to the storage system to perform administrative tasks if SSH is enabled for the storage system.

For storage systems shipped with Data ONTAP 8.0 or later, SecureAdmin is set up automatically and SSH is enabled by default. For systems running earlier releases of Data ONTAP, you must ensure that SecureAdmin has been set up and enabled before issuing SSH requests.

Data ONTAP provides 24 concurrent SSH administrative sessions. However, you can open only one SSH-interactive session at a time.

Step

1. From a UNIX client, enter the `ssh` command in one of the following formats:

```
ssh [-1|-2] username@{IP_addr|hostname} [command]
```

or

```
ssh [-1|-2] -l username {IP_addr|hostname} [command]
```

- The option `-1` forces SSH to use protocol version 1 only.

- The option `-2` forces SSH to use protocol version 2 only.
By default, SSH uses protocol version 2.
- `command` is not required for SSH-interactive sessions.

Examples of SSH requests

The following examples show how the user named "joe" that has been set up on the storage system named "mysystem" can issue a SSH request:

```
ssh joe@mysystem version
```

```
ssh joe@10.72.137.28 version
```

```
ssh -l joe 10.72.137.28 version
```

```
ssh -l joe@mysystem version
```

```
ssh -2 joe@mysystem version
```

Related concepts

[How to manage SSH](#) on page 53

[Rules that apply to console, Telnet, and SSH-interactive sessions](#) on page 43

Displaying the current SSH settings

If SSH has been enabled, you can use the `ssh` option to display the current SSH settings on your storage system.

Step

1. To display the current SSH settings, enter the following command at the storage system prompt:

```
options ssh
```

For more information about the SSH options and their default values, see the `na_options(1)` man page.

The current SSH settings on your storage system are displayed.

Example of `options ssh` output

```
mysystem> options ssh
ssh.access                *
ssh.enable                on
ssh.idle.timeout         600
ssh.passwd_auth.enable   on
ssh.port                  22
ssh.pubkey_auth.enable   on
ssh1.enable              on
```

```
ssh2.enable          on
mysystem>
```

The SSL protocol

The Secure Sockets Layer (SSL) protocol 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.

Data ONTAP supports SSLv2 and SSLv3. You should use SSLv3 because it offers better security protections than previous SSL versions.

As a precautionary measure due to security vulnerability CVE-2009-3555, the SSL renegotiation feature is disabled in Data ONTAP.

Next topics

[How to manage SSL](#) on page 61

[Setting up and starting SSL](#) on page 62

[Installing a certificate-authority-signed certificate](#) on page 63

[Testing certificates](#) on page 63

[Reinitializing SSL](#) on page 64

[Enabling or disabling SSL](#) on page 64

[Enabling or disabling SSLv2 or SSLv3](#) on page 65

How to manage SSL

SSL uses a certificate to provide a secure connection between the storage system and a Web browser. If your storage system does not have SSL enabled, you can set up SecureAdmin to enable SSL and allow administrative requests over HTTPS to succeed.

SecureAdmin is set up automatically on storage systems shipped with Data ONTAP 8.0 or later. For these systems, Secure protocols (including SSH, SSL, and HTTPS) are enabled by default, and nonsecure protocols (including RSH, Telnet, FTP, and HTTP) are disabled by default.

Two types of certificates are used—self-signed certificate and certificate-authority-signed certificate.

- 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.

Related concepts

[The default security settings](#) on page 50

Setting up and starting SSL

Setting up SSL enables Data ONTAP to generate a self-signed certificate.

Steps

1. Enter the following command at the storage system prompt:

```
secureadmin setup ssl
```

2. If SSL has been previously set up for the storage system, Data ONTAP asks you whether you want to continue.
 - Enter **Y** if you want to change the SSL setup.
 - Enter **N** to exit the SSL setup.
3. Enter information when Data ONTAP prompts you.

The information you are prompted to enter includes the following:

- Country, state, or province name
- Company or organization name
- Domain name
- Administrator email
- Days until expires
- Key length in bits

To use the default settings, press Enter at each of the prompts.

When the SSL setup is complete, Data ONTAP generates `secureadmin.pem` files and saves them in the appropriate subdirectories (`cert`, `key`, and `csr`) in the `/etc/keymgr` directory.

Related tasks

[Installing a certificate-authority-signed certificate](#) on page 63

[Testing certificates](#) on page 63

Installing a certificate-authority-signed certificate

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.

Steps

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
```

Related tasks

[Testing certificates](#) on page 63

Testing certificates

After installing either a self-signed certificate or a certificate-authority-signed certificate, you should test the certification to verify that it is installed correctly.

Steps

1. Start your Web browser.
2. Enter the following URL:

```
https://systemname/na_admin
```

systemname is the name of your storage system.

3. Click **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 might not have a secure connection.

Reinitializing SSL

You should reinitialize SSL if you change the domain name of the storage system. When you change the domain name of your system, the domain name recorded in the certificate becomes obsolete. As a result, the storage system is not authenticated after the domain name change, although the connection is still encrypted. 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.

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.

Steps

1. Disable SecureAdmin by entering the following command:

```
secureadmin disable ssl
```

2. Use the `secureadmin setup ssl` command to reinitialize SSL.

Related tasks

[Setting up and starting SSL](#) on page 62

Enabling or disabling SSL

Enabling SSL allows administrative requests over HTTPS to succeed. Disabling SSL disallows all administrative requests over HTTPS.

Before enabling SSL for the first time, you must set up SSL and install a certificate signed by a certificate authority.

Step

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

```
secureadmin {enable|disable} ssl
```

Use `enable` to start SSL. Use `disable` to deactivate SSL.

Related tasks

[Setting up and starting SSL](#) on page 62

[Installing a certificate-authority-signed certificate](#) on page 63

[Testing certificates](#) on page 63

Enabling or disabling SSLv2 or SSLv3

If your storage system has the SSL protocol enabled, you can specify the SSL version(s) to use.

Enabling the SSL versions alone does not enable the SSL protocol for the storage system. To use SSL, ensure that the protocol is enabled on your storage system.

SSLv3 is recommended over SSLv2 because SSLv3 offers better security protection than SSLv2. You can leave both SSL versions enabled, or you can disable one of them. In addition to enabling the SSL protocol, you must also have at least one SSL version enabled for the storage system to use SSL for communication.

Step

1. Enter the following command to enable or disable SSLv2 or SSLv3:

To enable or disable this SSL version:	Enter the following command:
SSLv2	<code>options ssl.v2.enable {on off}</code>
SSLv3	<code>options ssl.v3.enable {on off}</code>

Setting the option to `on` (the default) enables the SSL version on HTTPS and LDAP connections, if the following options are also set to `on`:

- `httpd.admin.ssl.enable` (for HTTPS)
- `ldap.ssl.enable` (for LDAP)

Setting the option to `off` disables the SSL version on HTTPS and LDAP connections.

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

For more information about LDAP, see the *Data ONTAP 7-Mode File Access and Protocols Management Guide*.

Related tasks

[Setting up and starting SSL](#) on page 62

Determining whether secure protocols are enabled

Data ONTAP displays information that shows whether secure protocols are enabled. The information helps you determine whether administrative transactions between the storage system and a client are being encrypted.

Step

1. Enter the following command:

```
secureadmin status
```

Information similar to the following is displayed:

```
ssh2    - active
ssh1    - inactive
ssl     - active
```

Related concepts

[The default security settings](#) on page 50

Enabling or disabling secure protocols

The `secureadmin` command allows you to enable or disable both SSH and SSL.

Step

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.

How to access a storage system by using Telnet

You can access a storage system from a client through a Telnet session if you enabled Telnet.

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`
- `telnet.distinct.enable`

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

Next topics

Starting a Telnet session on page 67

Terminating a Telnet session on page 68

Configuration for Telnet sessions on page 68

Related concepts

The default security settings on page 50

Rules that apply to console, Telnet, and SSH-interactive sessions on page 43

Starting a Telnet session

You start a Telnet session to connect to the storage system.

Before you begin

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`. 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 `na_options(1)` 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(8)` man pages.

About this task

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.

Steps

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.

The storage system responds with the password prompt.

4. Enter the password for the root or administrative user account.

Note: If no password is defined for the account, press Enter.

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.

Related concepts

Rules that apply to console, Telnet, and SSH-interactive sessions on page 43

Related tasks

Restricting protocol access on page 83

Terminating a Telnet session

You terminate a Telnet session to disconnect from the storage system.

Step

1. To log out of the storage system at the system prompt or at the console, do one of the following:

- Press Ctrl-] .
- Enter the following command:
`logout telnet`
- 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
```

Configuration for Telnet sessions

You can configure the Telnet sessions to display a banner message or specify the timeout period.

Next topics

Banner message configuration on page 69

Enabling or disabling the timeout period for Telnet or SSH-interactive sessions on page 69

Changing the timeout period for Telnet or SSH-interactive sessions on page 70

Banner message configuration

You can configure a banner message to appear at the beginning of a Telnet session to a storage system.

You 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 mysystem
Trying 192.0.2.132...
Connected to mysystem.xyz.com
Escape character is '^]'.

This system is for demonstrations only.

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

Enabling or disabling the timeout period for Telnet or SSH-interactive sessions

You can enable or disable the timeout period for Telnet or SSH-interactive sessions. If the timeout period is enabled, Telnet or SSH-interactive connections are automatically disconnected after the number of minutes specified by the `autologout.telnet.timeout` option has elapsed.

Step

1. To enable or disable the timeout period for Telnet or SSH-interactive sessions, enter the following command:

```
options autologout.telnet.enable [on|off]
```

The default is `on`, which causes Telnet or SSH-interactive connections to be disconnected automatically after the number of minutes specified by the `autologout.telnet.timeout` option has elapsed.

Any change to the `autologout.telnet.enable` option requires a logout before it takes effect.

Changing the timeout period for Telnet or SSH-interactive sessions

You can change the timeout period for Telnet or SSH-interactive sessions. By default, Telnet and SSH-interactive sessions have a timeout period of 60 minutes.

Ensure that the `autologout.telnet.enable` option is set to on for the `autologout.telnet.timeout` option to take effect.

Step

1. To change the timeout period for Telnet or SSH-interactive sessions, 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. The default is 60 minutes.

How to access a storage system by using a Remote Shell connection

If the `rsh.enable` option is set to on, you can access a storage system to perform administrative tasks by using a Remote Shell (RSH) connection.

You can access a storage system by using an RSH connection with a trusted remote host that is listed in the `/etc/hosts.equiv` file on the root volume.

You can also use a user name and a password to establish an RSH connection from an administration host that is not listed in the `/etc/hosts.equiv` file. However, 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 plain text over the network. Any program that captures network traffic when the password is sent will record the password. To avoid exposing the password when you issue RSH commands, it is best to log in as root on a client listed in the storage system's `/etc/hosts.equiv` file.

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.

Next topics

[When to use RSH commands with user names and passwords](#) on page 71

[Accessing a storage system from a UNIX client by using RSH](#) on page 72

Accessing a storage system from a Windows client by using a Remote Shell application on page 73

Commands not accepted when using RSH on page 73

How to reset options to default values from RSH on page 74

Displaying RSH session information on page 74

Related concepts

The default security settings on page 50

How to specify administration hosts on page 80

Public-key-based authentication on page 56

Related tasks

Restricting protocol access on page 83

Adding administration hosts on page 81

Removing administration hosts on page 81

Restricting protocol access on page 83

When to use RSH commands with user names and passwords

Depending on the UNIX host you use and how you log in to the UNIX host, you might need to supply a user name and a password when using the RSH protocol to run a command on the storage system.

If the UNIX host you use is not listed in the storage system's `/etc/hosts.equiv` file, you must supply both a user name and a password when using the RSH protocol to run a command on the storage system.

If the UNIX host you use is listed in the storage system's `/etc/hosts.equiv` file and you are logged in as root on the UNIX host, you do not need to supply a user name or a password when using the RSH protocol to run a command on the storage system.

If the UNIX host you use is listed in the storage system's `/etc/hosts.equiv` file and you are logged in as a user other than root on the UNIX host, the following rules apply when using the RSH protocol to run a command on the storage system:

- If the user name is listed with the host name in the `/etc/hosts.equiv` file, supplying a user name is optional. You do not need to supply a password.
- If the user name is not listed with the host name in the `/etc/hosts.equiv` file, you must supply both a user name and a password.

The user name can be root or the name of an administrative user that is 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.

Accessing a storage system from a UNIX client by using RSH

You can use an RSH connection to access a storage system from a UNIX client to perform administrative tasks.

Before you begin

Ensure that the `rsh.enable` option is set to on.

Step

1. Do one of the following:

- If the UNIX host name or the user name you use is not specified in the `/etc/hosts.equiv` file on the root volume of the storage system, enter the `rsh` command in the following format:

```
rsh hostname_or_ip -l username:password command
```

- If the UNIX host name and the user name you use are specified in the `/etc/hosts.equiv` file on the root volume of the storage system, enter the `rsh` command in the following format:

```
rsh hostname_or_ip [-l username] command
```

`hostname_or_ip` is the host name or IP address of the storage system.

Note: You can also specify the IP address by using the `rsh.access` option.

`command` is the Data ONTAP command you want to run over the RSH connection.

Examples of RSH requests

The following `rsh` command uses a user name, `carl`, and a password, `mypass`, to access the storage system, `myfiler`, to run the Data ONTAP `version` command:

```
rsh myfiler -l carl:mypass version
```

The following `rsh` command uses a user name, `carl`, and a password, `mypass`, to access the storage system whose IP address is `192.0.2.66` to run the Data ONTAP `version` command:

```
rsh 192.0.2.66 -l carl:mypass version
```

The following `rsh` command runs the Data ONTAP `version` command from a UNIX host that is specified in the `/etc/hosts.equiv` file of the storage system, `myfiler`:

```
rsh myfiler version
```

Related tasks

[Restricting protocol access](#) on page 83

Accessing a storage system from a Windows client by using a Remote Shell application

You can use a Remote Shell application to access a storage system from a Windows client to perform administrative tasks.

Before you begin

Ensure that the `rsh.enable` option is set to `on`.

Ensure that the Windows client you use is a trusted host specified in the `/etc/hosts.equiv` file on the root volume of the storage system.

Steps

1. Run the Remote Shell application on the Windows client.
2. From the the Remote Shell application, enter the `rsh` command in the following format:

```
rsh hostname_or_ip [-l username:password] command
```

`hostname_or_ip` is the host name or IP address of the storage system.

Note: You can also specify the IP address by using the `rsh.access` option.

`command` is the Data ONTAP command you want to run over the RSH connection.

Examples of RSH requests

The following `rsh` command uses a user name, `carl`, and a password, `mypass`, to access the storage system, `myfiler`, to run the Data ONTAP `version` command:

```
rsh myfiler -l carl:mypass version
```

The following `rsh` command uses a user name, `carl`, and a password, `mypass`, to access the storage system whose IP address is `192.0.2.66` to run the Data ONTAP `version` command:

```
rsh 192.0.2.66 -l carl:mypass version
```

Related tasks

[Restricting protocol access](#) on page 83

Commands not accepted when using RSH

You cannot execute several commands when you use RSH.

The commands that you cannot execute when you use RSH include the following:

- `arp`

- `orouted`
- `ping`
- `routed`
- `savecore`
- `setup`
- `traceroute`

How to reset 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, 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 "\""
```

Displaying RSH session information

The `rshstat` command displays information about RSH sessions, such as the number of RSH sessions invoked, the number of currently active RSH sessions, and the highest number of concurrently active RSH sessions.

Step

1. Enter the following command:

```
rshstat [ -a | -t ]
```

Without any options, `rshstat` displays the following information:

- The number of RSH sessions invoked since booting the storage system
- The number of currently active RSH sessions
- The highest number of concurrently active RSH sessions since booting the storage system
- The maximum concurrent RSH sessions allowed

The `-a` option displays the following additional information:

- The RSH session number
- The command the RSH session is executing

Note: `rsh shell` in the command field means that the RSH session is being initiated.

- The remote client's IP address for the RSH session
- The last string written into the audit log for the RSH session

The `-t` option displays the amount of time the command is running in milliseconds, in addition to the information displayed by the `-a` option. The time information includes:

- The total time used for running the command
- The protocol connection time
- The host lookup (gethost) information time

Example

```
toaster> rshstat
Session Invocations: 9
Current Active Sessions: 2
Active High Sessions: 3
Maximum Available Sessions: 24

toaster> rshstat -a
Session Invocations: 9
Current Active Sessions: 2
Active High Sessions: 3
Maximum Available Sessions: 24

0: sysstat [from 192.0.2.66] (50% 0 0 0 178 219 0 0 0 0 >60 )
-----
1: nfsstat [from 2001:0DB8:85A3:0:0:8A2E:0370:99] (0 0 0 0 0 0 0 0)
-----

toaster> rshstat -t
Session Invocations: 9
Current Active Sessions: 2
Active High Sessions: 3
Maximum Available Sessions: 24

0: sysstat [from 192.0.2.66] (50% 0 0 0 178 219 0 0 0 0 >60 )
Command Time: 123ms
Connection Time: 123ms
Gethost Time: 123ms
-----
1: nfsstat [from 2001:0DB8:85A3:0:0:8A2E:0370:99] (0 0 0 0 0 0 0 0)
Command Time: 3490ms
Connection Time: 3490ms
Gethost Time: 3490ms
```

How to access a storage system by using FilerView

You can use FilerView to access a storage system. FilerView is a 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 or an RSH session, 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 Help explains Data ONTAP features and how to use them.

FilerView supports the following browsers:

- Microsoft Internet Explorer® 6 and 7
- Mozilla® Firefox® 2.0
- Mozilla Suite 1.7 or later

The following options control access to FilerView:

- `httpd.admin.access`
Restricts HTTP access to FilerView. If this value is set, `trusted.hosts` is ignored for FilerView access.
- `httpd.admin.enable`
Enables HTTP access to FilerView.
- `httpd.admin.ssl.enable`
Enables HTTPS access to FilerView.
- `httpd.admin.top-page.authentication`
Specifies whether the top-level FilerView administration Web page prompts for user authentication.

Note: On storage systems shipped with Data ONTAP 8.0 or later, secure protocols (including SSH, SSL, and HTTPS) are enabled by default, and nonsecure protocols (including RSH, Telnet, FTP, and HTTP) are disabled by default.

For information about how to use these options, see the `na_options(1)` man pages.

For information about using the `secureadmin` command to set up Secure Sockets Layer (SSL), see the `na_secureadmin(1)` man page.

Next topics

[Accessing a storage system from a client by using FilerView](#) on page 76

[The FilerView interface](#) on page 78

[Read-only access to FilerView](#) on page 79

Related concepts

[The default security settings](#) on page 50

Accessing a storage system from a client by using FilerView

You can use FilerView to manage most storage system functions and view information about the storage system.

Before you begin

The browser must have Java and JavaScript enabled.

If your version of Microsoft Windows does not include Java support, you must download the Java Runtime Environment (JRE) separately to ensure that FilerView functions properly.

To use FilerView over HTTP, ensure that the `httpd.admin.enable` option is set to `on`.

To use FilerView over HTTPS, ensure that the `httpd.admin.ssl.enable` option is set to `on`.

Steps

1. Start your Web browser.
2. Enter the FilerView URL in one of the following formats:

To use...	Enter...
HTTP	<code>http://filer_name_or_IP/na_admin</code>
HTTPS	<code>https://filer_name_or_IP/na_admin</code>

filer_name_or_IP can be one of the following:

- The short name of the storage system
- The fully qualified name of the storage system
- The IPv4 address of the storage system

Using the HTTPS format allows you to access FilerView securely. For storage systems shipped with Data ONTAP 8.0 or later, secure protocols (including SSH, SSL, and HTTPS) are enabled by default, and nonsecure protocols (including RSH, Telnet, FTP, and HTTP) are disabled by default. If your storage system currently does not have secure protocols enabled, you can use the `secureadmin` command to enable SSL.

3. If the `httpd.admin.top-page.authentication` option is set to `on` (the default), a login dialog box appears, prompting you for user authentication before you can access the top-level FilerView administration Web page. Enter a user name and password. Then click **OK**.

Note: If the `httpd.admin.top-page.authentication` option is set to `off`, the top-level FilerView administration Web page appears without user authentication.

4. Click **FilerView**.

- If the storage system is password protected, you are prompted for a user name and password.
- If the storage system is not password protected, FilerView is launched, and a window appears with a list of categories in the left pane and the System Status information in the right pane.

Examples of FilerView URLs

The following FilerView URL uses the storage system name (`mssystem`) and HTTP to access the storage system:

```
http://mssystem/na_admin
```

The following FilerView URL uses the fully qualified name of the storage system (mysystem.mycompany.com) and HTTPS to access the storage system:

```
https://mysystem.mycompany.com/na_admin
```

The following FilerView URL uses the IPv4 address of the storage system (192.0.2.66) and HTTP to access the storage system:

```
http://192.0.2.66/na_admin
```

Related concepts

[How to manage SSL](#) on page 61

[The default security settings](#) on page 50

The FilerView interface

The FilerView interface consists of the following main elements: the left pane, the right pane, the title pane, and the Help buttons.

- Left pane

The left pane contains an expandable list of topics:

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

If you select the Manage, Configure, or Report functions from the left pane, the right pane 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.
- Title pane

The title pane contains the name of the function you select from the left pane, followed by the path to the function. For example, if you select Report in the Volumes category, the title pane shows the path as **Volumes > Report**.
- Help buttons

Help buttons are indicated with a question mark (?) and are situated next to the categories in the left pane and in the title pane. 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.

When you click the Help button next to a category, a two-pane Help window appears. The left pane displays an expandable table of contents, with additional tabs at the top labeled Index and Search.

Read-only access to FilerView

Users with the `filerview-readonly` capability have read-only access to FilerView.

Users with read-only FilerView access can view objects on the storage system that FilerView manages. They are not allowed to add or modify objects through FilerView.

Related concepts

[Supported capability types](#) on page 129

How to manage access from administration hosts

An administration host can be any workstation that is either an NFS or a CIFS client on the network.

Next topics

[Reasons to designate a workstation as an administrative host](#) on page 79

[Administration host privileges](#) on page 80

[Requirements for using a client](#) on page 80

[How to specify administration hosts](#) on page 80

[Adding administration hosts](#) on page 81

[Removing administration hosts](#) on page 81

Reasons to designate a workstation as an administrative host

You designate a workstation as an administration host to limit access to the storage system's root file system, to provide a text editor to edit configuration files, or to provide the ability to administer a storage system remotely.

During the setup process, you are prompted to designate a workstation on the network as an administration host. For more information about the setup process, see the *Data ONTAP 7-Mode Software Setup Guide*.

When you designate a workstation as an administration host, the storage system's root file system (`/vol/vol0` by default) is accessible only to the specified workstation in the following ways:

- As a share named C\$, if the storage system is licensed for the CIFS protocol
- By NFS mounting, if the storage system is licensed for the NFS protocol

If you do not designate a workstation as an administration host, the storage system's root file systems are available to all workstations on the network. As a result, any user can gain access to the storage system's root file system and change or remove storage system configuration files in the `/etc` directory.

You can designate additional administration hosts after setup by modifying the storage system's NFS exports and CIFS shares.

Administration host privileges

After the setup procedure is completed, the storage system grants root permissions to the administration host.

If the administration host you use is an NFS client, you have the privilege to perform the following tasks:

- Mount the storage system root directory and edit configuration files from the administration host.
- Enter Data ONTAP commands by using an RSH connection (if RSH is enabled on the storage system) or an SSH connection (if SSH is enabled on the storage system).

If the administration host you use is a CIFS client, you have the privilege to edit configuration files from any CIFS client as long as you connect to the storage system as root or Administrator.

Requirements for using a client

An NFS or CIFS client must meet the requirements to manage the storage system.

If you plan to use an NFS client to manage the storage system, the NFS client must meet the following requirements:

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

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

How to specify administration hosts

Administration hosts are specified in the `/etc/hosts.equiv` file.

You use one of the following formats to specify an administration host:

- `hostname_or_ip [username]` or `hostname_or_ip ["user name"]` for a user on a host
- `+@netgroup [username]` for a group of hosts

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.

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 specify each user's entry for a single host using `hostname_or_ip [username]`. You can

also specify a group of hosts using `+`*@netgroup* [*username*] to allow a particular user to access the storage system from a group of hosts.

- If *hostname_or_ip* specifies an NFS client, or if `+`*@netgroup* specifies a group of NFS hosts, the user name is optional. If you do not specify a user name, you must be the root user on that NFS client or the root user on the host in the host group to execute a Data ONTAP command through a Remote Shell connection.
- If *hostname_or_ip* specifies a CIFS client, you must enter the user name for that CIFS client.

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

```
nfsclient1
client1 carl
client1 peter
client2 lena
client2 root
client3 fred
client3 root
+@sysadmins joe smith
```

For more information, see the `na_hosts.equiv(5)` man page.

Adding administration hosts

You can designate additional NFS clients or CIFS clients as administration hosts by editing the `/etc/hosts.equiv` file.

Steps

1. Open the `/etc/hosts.equiv` configuration file with an editor.
2. Add the group of hosts or 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

You can remove an NFS client or CIFS client from the administration hosts list by editing the `/etc/hosts.equiv` file.

Steps

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

Methods for controlling storage system access

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

Next topics

Controlling Telnet access using host names on page 82

Restricting protocol access on page 83

Controlling mount privilege on page 84

Controlling file ownership change privileges on page 84

Controlling anonymous CIFS share lookups on page 85

Options that help maintain security on page 85

Allowing only secure access to the storage system on page 87

Related concepts

The default security settings on page 50

Controlling Telnet access using host names

You can disable Telnet access for all hosts, restrict Telnet access to up to five hosts, or allow Telnet access for all hosts.

Steps

1. Access the storage system command line through the console or through a Telnet session.
2. Do one of the following:

If...	Then...
You want to disable Telnet access for all hosts	Enter the following command: <code>options trusted.hosts -</code>
You want to restrict Telnet access to up to five hosts	Enter the following command: <code>options trusted.hosts host1[, ..., host5]</code>
You want to allow Telnet access for all hosts	Enter the following command: <code>options trusted.hosts *</code>

Restricting protocol access

If a protocol is enabled for Data ONTAP, you can restrict the protocol's access to the storage system by specifying the host name, IP address, or network interface name.

Step

1. At the storage system prompt, enter one of the following commands:

If you want to restrict a protocol's access to the storage system by using... **Enter...**

host name or IP address	<code>options protocol.access host=[hostname IP_address]</code>
-------------------------	---

network interface name	<code>options protocol.access if=interface_name</code>
------------------------	--

- *protocol* is the name of the protocol you want to allow access to the storage system. It can be `rsh`, `telnet`, `ssh`, `httpd`, `httpd.admin`, `snmp`, `ndmpd`, `snapmirror`, or `snapvault`.
- *hostname* is the name of the host to which you want to allow access by using *protocol*.
- *IP_address* is the IP address of the host to which you want to allow access by using *protocol*.
- *interface_name* is the network interface name of the host to which you want to allow access by using *protocol*.

Note: If the `telnet.access` option is not set to `legacy`, the `trusted.hosts` option is ignored for Telnet. If the `httpd.admin.access` option is not set to `legacy`, the `trusted.hosts` option is ignored for `httpd.admin`. If the `snapmirror.access` option is not set to `legacy`, the `/etc/snapmirror.allow` file is ignored for SnapMirror destination checking.

For more information about controlling protocol access to a storage system by using multiple host names, IP addresses, and network interfaces, see the `na_protocolaccess(8)` man page.

For information about SNMP, see the *Data ONTAP 7-Mode Network Management Guide*.

For information about NDMP, see the *Data ONTAP 7-Mode Data Protection Tape Backup and Recovery Guide*.

For information about SnapMirror or SnapVault, see the *Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide*.

Related tasks

[Allowing only secure access to the storage system](#) on page 87

Controlling mount privilege

You can control the NFS mount privilege for the storage system's volumes by restricting the mount privilege to only the root user using privileged ports.

About this task

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.

Steps

1. Access the storage system command line through the console or through a Telnet session.
2. Do one of the following:

If you want to ...	Enter the following command ...
Restrict the mount privilege to only the root user using privileged ports (ports 1 through 1,024)	<code>options nfs.mount_rootonly on</code>
Allow the mount privilege for all users on all ports	<code>options nfs.mount_rootonly off</code>

Controlling file ownership change privileges

You can control who has privileges to change directory and file ownership.

About this task

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.

Steps

1. Access the storage system command line through the console or through a Telnet session.
2. Do one of the following:

If...	Then...
You want to restrict the privilege of changing directory and file ownership to the root user	Enter the following command: <code>options waf1.root_only_chown on</code>

If...	Then...
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

You can control whether anonymous CIFS users can look up CIFS shares, users, or groups on a storage system.

Steps

1. Access the storage system command line through the console or through a Telnet session.
2. Do one of the following:

If ...	Enter the following command ...
You do not want to set access restrictions for anonymous share lookups	options cifs.restrict_anonymous 0
You do not want to allow enumeration of users and shares	options cifs.restrict_anonymous 1
You want to fully restrict anonymous share lookups	options cifs.restrict_anonymous 2

The default value for the `cifs.restrict_anonymous` option is 0. The restrictions do not apply to mapped null users. For more information, see the `na_options(1)` man page.

Options that help maintain security

Several options are available to help you maintain storage system security.

The following table shows the options that help maintain security.

Option	Description
<code>trusted.hosts</code>	<p>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.</p> <p>This value is ignored for Telnet access if the <code>telnet.access</code> option is set. It is also ignored for administrative HTTP access if the <code>httpd.admin.access</code> option is set.</p>

Option	Description
telnet.access	<p>Controls which hosts can access the storage system through a Telnet session for administrative purposes.</p> <p>You can restrict Telnet access to the storage system by specifying host names, IP addresses, or network interface names. If this value is set, the <code>trusted.hosts</code> option is ignored for Telnet.</p>
telnet.distinct.enable	<p>Controls whether the Telnet and the SSH environments are shared with or separate from the console environment.</p> <p>When the option is set to <code>off</code>, a Telnet or an SSH session is shared with a console session. A Telnet or an SSH user and a console user can view each other's inputs or outputs, and they acquire the privileges of the last Telnet, SSH, or console user who logged in.</p> <p>You can keep the Telnet and the SSH environments separate from the console environment by ensuring that the option is set to <code>on</code>.</p> <p>If the setting for this option is changed during a Telnet or an SSH session, the change does not go into effect until the next Telnet or SSH login.</p>
rsh.access	<p>Controls which hosts can access the storage system through a Remote Shell session for administrative purposes.</p> <p>You can restrict Remote Shell access to the storage system by specifying host names, IP addresses, or network interface names.</p>
ssh.access	<p>Controls which hosts can access the storage system through a Secure Shell session for administrative purposes.</p> <p>You can restrict Secure Shell access to the storage system by specifying host names, IP addresses, or network interface names.</p>
nfs.mount_rootonly	<p>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.</p> <p>This option is applicable only if the NFS protocol is licensed.</p>
wafl.root_only_chown	<p>Controls whether all users or only the root user can change directory and file ownership.</p> <p>This option is applicable only if the NFS protocol is licensed.</p>
cifs.restrict_anonymous	<p>Controls whether anonymous CIFS users can look up CIFS shares, users, or groups on a storage system.</p> <p>This option is applicable only if the CIFS protocol is licensed.</p>

For more information about the options in this table, see the `na_options(1)` and the `na_protocolaccess(8)` man pages.

Related tasks

[Restricting protocol access](#) on page 83

Allowing only secure access to the storage system

If you want to allow only secure access to your storage system, enable secure protocols and disable nonsecure protocols. You should also set password rule options to enhance password security.

About this task

On storage systems shipped with Data ONTAP 8.0 or later, secure protocols (including SSH, SSL, and HTTPS) are enabled and nonsecure protocols (including RSH, Telnet, FTP, and HTTP) are disabled by default.

Steps

1. Use the `secureadmin` commands to set up and enable the secure protocols, SSH and SSL.

After you have set up SecureAdmin to enable SSH and SSL, the following options are set to on:

- `options ssh.enable`
- `options ssh2.enable` (if you enabled SSHv2 during SecureAdmin setup)
- `options ssh.passwd_auth.enable`
- `options ssh.pubkey_auth.enable`
- `options httpd.admin.ssl.enable`

2. Disable nonsecure protocols.

To disable the following access to the storage system...	Enter the following at the storage system prompt...
RSH	<code>options rsh.enable off</code>
Telnet	<code>options telnet.enable off</code>
FTP	<code>options ftpd.enable off</code>
HTTP	<code>options httpd.enable off</code>

Note: This option controls HTTP access to the storage system.

To disable the following access to the storage system...	Enter the following at the storage system prompt...
HTTP access to FilerView	<p>options httpd.admin.enable off</p> <p>Note: This option controls HTTP access to FilerView.</p> <p>To use HTTPS to access FilerView securely, ensure that the <code>httpd.admin.ssl.enable</code> option is set to <code>on</code>. If both the <code>httpd.admin.enable</code> option and the <code>httpd.admin.ssl.enable</code> option are set to <code>off</code>, you cannot access the storage system by using FilerView.</p>
SSHv1	<p>options ssh1.enable off</p> <p>Note: Ensure that the <code>ssh.enable</code> option and the <code>ssh2.enable</code> option are set to <code>on</code>.</p>

3. Ensure that the following password options are set:

- `options security.passwd.rules.everyone on`
This option ensures that password composition is checked for all users, including root and Administrator.
- `options security.passwd.rules.history 6`
This option prevents users from reusing any of the six previously used passwords.

Related concepts

[Secure protocols and storage system access](#) on page 49

[The default security settings](#) on page 50

[How to access a storage system by using FilerView](#) on page 75

[Options that manage password rules](#) on page 143

How to manage the root volume

The storage system's root volume contains special directories and configuration files that help you administer your storage system.

The root volume is installed at the factory on FAS systems and on V-Series systems ordered with disk shelves.

Note: For a V-Series system that does not have a disk shelf, you need to install the root volume on the third-party storage. If you use a FlexVol volume for the root volume, you must ensure that it has a space guarantee of `volume`. For more information, see the *Data ONTAP 7-Mode Software Setup Guide*.

Unless the installer selected a unique volume name during setup, the default root volume name, `/vol/vol0`, is used.

The root aggregate contains the root volume. Your storage system is shipped with the root volume in a 32-bit aggregate. You can designate a different volume to be the new root volume. Starting in Data ONTAP 8.0.1, you can use a 64-bit volume for the root volume.

By default, the storage system is set up to use a hard disk drive (HDD) aggregate for the root aggregate. When no HDDs are available, the system is set up to use a solid-state drive (SSD) aggregate for the root aggregate. If you want to change the root aggregate, you can choose either an HDD aggregate or an SSD aggregate to be the root aggregate (by using `aggr options aggr_name root`), provided that the corresponding type of disk drives are available on the system.

For more information about traditional and FlexVol volumes, see the *Data ONTAP 7-Mode Storage Management Guide*.

Next topics

[Recommendations regarding the root volume](#) on page 89

[Size requirement for root FlexVol volumes](#) on page 91

[Default directories in the root volume](#) on page 92

[How to access the default directories on the storage system](#) on page 97

[Changing the root volume](#) on page 100

Recommendations regarding the root volume

There are recommendations and considerations to keep in mind when choosing what kind of volume to use for the root volume.

The following are the general recommendations regarding the root volume:

- Root volumes can use either FlexVol or traditional volumes.

Note: Data ONTAP 8.0 or later allows you to create only a new FlexVol root volume, not a new traditional root volume, from the boot menu. However, preexisting traditional root volumes are still supported.

- For small storage systems where cost concerns outweigh resiliency, a FlexVol based root volume on a regular aggregate might be more appropriate.
- Avoid storing user data in the root volume, regardless of the type of volume used for the root volume.
- For a V-Series system with a disk shelf, the root volume can reside on the disk shelf (recommended) or on the third-party storage. For a V-Series system that does not have a disk shelf, the root volume resides on the third-party storage. You can install only one root volume per V-Series system, regardless of the number of storage arrays or disk shelves that the V-Series system uses for storage.

The following are additional facts and considerations if the root volume is on a disk shelf:

- Data ONTAP supports two levels of RAID protection, RAID4 and RAID-DP. RAID4 requires a minimum of two disks and can protect against single-disk failures. RAID-DP requires a minimum of three disks and can protect against double-disk failures. The root volume can exist as the traditional stand-alone two-disk volume (RAID4) or three-disk volume (RAID-DP). Alternatively, the root volume can exist as a FlexVol volume that is part of a larger hosting aggregate.
- Smaller stand-alone root volumes offer fault isolation from general application storage. On the other hand, FlexVol volumes have less impact on overall storage utilization, because they do not require two or three disks to be dedicated to the root volume and its small storage requirements.
- If a FlexVol volume is used for the root volume, file system consistency checks and recovery operations could take longer to finish than with the two- or three-disk traditional root volume. FlexVol recovery commands work at the aggregate level, so all of the aggregate's disks are targeted by the operation. One way to mitigate this effect is to use a smaller aggregate with only a few disks to house the FlexVol volume containing the root volume.
- In practice, having the root volume on a FlexVol volume makes a bigger difference with smaller capacity storage systems than with very large ones, in which dedicating two disks for the root volume has little impact.
- For higher resiliency, use a separate two-disk root volume.

Note: You should convert a two-disk root volume to a RAID-DP volume when performing a disk firmware update, because RAID-DP is required for disk firmware updates to be nondisruptive. When all disk firmware and Data ONTAP updates have been completed, you can convert the root volume back to RAID4.

For Data ONTAP 7.3 and later, the default RAID type for traditional root volume is RAID-DP. If you want to use RAID4 as the raid type for your traditional root volume to minimize the number of disks required, you can change the RAID type from RAID-DP to RAID4 by using `vol options vol0 raidtype raid4`.

The following requirement applies if the root volume is on a storage array:

- For storage systems whose root volume is on a storage array, only one array LUN is required for the root volume regardless of whether the root volume is a traditional volume or a FlexVol volume.

Related concepts

[Size requirement for root FlexVol volumes](#) on page 91

Related tasks

[Changing the root volume](#) on page 100

Size requirement for root FlexVol volumes

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.

It is possible to create a FlexVol volume that is too small to be used as the root volume. 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. Data ONTAP also prevents you from resizing the root volume below the minimum allowed size or changing the space guarantee for the root volume.

The minimum size for a root FlexVol volume depends on your storage system model. The following table lists the required minimum size for root volumes. Check to ensure that the FlexVol volume to be used as the root volume meets the minimum size requirement. If you are using third-party storage, ensure that the array LUN you are using for the root volume is large enough to meet the minimum size requirement for the root volume.

Storage system model	Minimum root FlexVol volume size
FAS2040	160 GB
3040	160 GB
3070	230 GB
3140	160 GB
3160	240 GB
3170	250 GB
3210	100 GB
3240	150 GB
3270	300 GB
6030	250 GB

Storage system model	Minimum root FlexVol volume size
6040	250 GB
6070	250 GB
6080	250 GB
6210	300 GB
6240	300 GB
6280	300 GB
SA320	300 GB
SA300	230 GB
SA600	250 GB

Note: You cannot increase the root volume to more than 95 percent of the available aggregate size. The output of `df -A` displays the space used by the aggregates in the system.

The minimum array LUN size shown in the *V-Series Support Matrix* does not apply to the root volume.

Default directories in the root volume

The root volume contains the `/etc` directory and the `/home` directory, which were created when the storage system was set up. The `/etc` directory contains configuration files that the storage system needs in order to operate. The `/home` directory is a default location you can use to store data.

For a V-Series system that has a disk shelf, the root volume can reside on the disk shelf (recommended) or on the third-party storage. For a V-Series system that does not have a disk shelf, the root volume resides on the third-party storage. Regardless of how many third-party storage arrays are behind the V-Series system, each V-Series system can have only one root volume.

Next topics

[Permissions for the default directories](#) on page 92

[The `/etc` directory](#) on page 93

Permissions for the default directories

Permissions are assigned to the default directories when `setup` finishes.

The following table shows the permissions.

This directory...	From this client...	Has these permissions
The <code>/etc</code> directory	NFS	<ul style="list-style-type: none"> • Full permissions for the root user on the administration host (<code>-rwx</code>) • No permissions for any other user or host
	CIFS	<ul style="list-style-type: none"> • 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 <code>/home</code> directory	NFS	Permissions associated with individual users and with groups through a UNIX security database
	CIFS	Permissions for the <code>HOME\$</code> share are Full Control for Everyone

The `/etc` directory

The `/etc` directory is contained in the root directory. It 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.

Next topics

[The configuration files](#) on page 93

[How you edit configuration files](#) on page 94

[Enabling an NFS client to edit configuration files](#) on page 95

[Editing configuration files from a CIFS client](#) on page 96

[The `/etc/messages` file](#) on page 96

[The `/etc/usermap.cfg` file and the `/etc/quotas` file](#) on page 96

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 listed 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 quotas file, see the *Data ONTAP 7-Mode Storage Management Guide*. For more information about other editable configuration files, see the man pages.

Related concepts

[Startup configuration for the storage system](#) on page 165

How you edit configuration files

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.

Data ONTAP requires that the following configuration files be terminated with a carriage return. When you edit these files, be sure to insert a carriage return after the last entry:

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

Attention: 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
- cluster_config/*
- lclgroups.cfg
- filesid.cfg
- sysconfigtab
- registry.*

The following table provides the hard limits for some of the configuration files in the /etc directory.

File name	Limits
/etc/exports	Maximum entry size of 4,096 characters. Maximum number of entries are 10,240.
/etc/group	Maximum line size of 256 characters. No file size limit.
/etc/hosts	Maximum line size is 1,022 characters. Maximum number of aliases is 34. No file size limit.

File name	Limits
/etc/netgroup	<p>Maximum entry size of 4,096 characters.</p> <p>Maximum netgroup nesting limit is 1,000.</p> <p>No file size limit.</p> <p>Netgroup lookup is case-sensitive and must match the case used by DNS or NIS servers for host lookup.</p>
/etc/passwd	<p>Maximum line size of 256 characters.</p> <p>No file size limit.</p>
/etc/resolv.conf	<p>Maximum line size is 256.</p> <p>Maximum number of name servers is 3.</p> <p>Maximum domain name length is 256.</p> <p>Maximum search domains limit is 6.</p> <p>Total number of characters for all search domains is limited to 256.</p> <p>No file size limit.</p>

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.

The following steps to authorize access to the root file system are intended for an NFS client that is not specified as an administration host.

Steps

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

You can use a CIFS client to access the storage system's C\$ share and select a file to edit.

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.

Steps

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.

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.

The `/etc/messages` file

By default, all system messages of level INFO and higher are sent to the console and to the `/etc/messages` file, which enables you to see a record of events on your storage system and use scripts to parse for particular events.

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 page.

Related concepts

[Message logging](#) on page 160

[How to access the default directories on the storage system](#) on page 97

Related tasks

[Accessing log files using HTTP or HTTPS](#) on page 100

The `/etc/usermap.cfg` file and the `/etc/quotas` file

The `/etc/usermap.cfg` file is used by Data ONTAP to map user names. The `/etc/quotas` file consists of entries to specify a default or explicit space or file quota limit for a qtree, group, or user.

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.

How to access the default directories on the storage system

You can access the default directories from an NFS client, a CIFS client, or with FTP. You can also access your log files by using HTTP or HTTPS.

Next topics

[Accessing the /etc directory from an NFS client](#) on page 97

[Accessing the /etc directory from a CIFS client](#) on page 97

[Accessing the /etc directory with FTP](#) on page 98

[Accessing the /home directory from an NFS client](#) on page 98

[Accessing the /home directory from a CIFS client](#) on page 99

[Accessing the /home directory with FTP](#) on page 99

[Accessing log files using HTTP or HTTPS](#) on page 100

Accessing the /etc directory from an NFS client

You can access the `/etc` directory from an NFS client to manage your storage system.

Steps

1. Mount the following path:

```
filer: /vol/vol0
```

filer is the name of your storage system.

You now have access to the storage system's root directory.

2. Change directories to the storage system's `/etc` directory by using the following command:

```
cd mountpoint/etc
```

mountpoint is the name of the storage system's mountpoint on the NFS client.

Accessing the /etc directory from a CIFS client

You can access the `/etc` directory from a CIFS client to manage your storage system.

Steps

1. Map a drive to the following path:

```
\\filer\C$
```

filer is the name of your storage system.

You have access to the storage system root directory.

2. Double-click the `/etc` folder to access the content.

Accessing the `/etc` directory with FTP

You can use the File Transfer Protocol (FTP) to access the `/etc` directory of your storage system.

Steps

1. Enable FTP access on the storage system by entering the following command:

```
options ftpd.enable on
```

2. Set the default home directory to `/etc` by entering the following command:

```
options ftpd.dir.override /vol/vol0/etc
```

For more information about FTP, see the *Data ONTAP 7-Mode File Access and Protocols Management Guide* and the `na_options(1)` man page.

3. Connect to the storage system from a client by using FTP.
4. Use the FTP `get` command to copy files from the storage system to your client so you can edit them.
5. Use the FTP `put` command to copy the edited files from your client to the storage system.

Related concepts

[The default security settings](#) on page 50

Accessing the `/home` directory from an NFS client

You can access the `/home` directory of your storage system from an NFS client to manage the storage system.

Step

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

You can access the /home directory of your storage system from a CIFS client to manage the storage system.

Step

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 the /home directory with FTP

You can use FTP to access the /home directory of your storage system.

Steps

1. Enable FTP access on the storage system by entering the following command:

```
options ftpd.enable on
```

2. Set the default home directory by entering the following command:

```
options ftpd.dir.override /vol/vol0/home
```

For more information about FTP, see the *Data ONTAP 7-Mode File Access and Protocols Management Guide* and the `na_options(1)` man page.

3. Connect to the storage system from a client by using FTP.
4. Use the FTP `get` command to copy files from the storage system to your client so you can edit them.
5. Use the FTP `put` command to copy the edited files from your client to the storage system.

Related concepts

[The default security settings](#) on page 50

Accessing log files using HTTP or HTTPS

You can access your log files by using HTTP or HTTPS, whichever is enabled for your storage system.

Before you begin

Ensure that the `httpd.autoindex.enable` option is set to `on` and that the `httpd.admin.access` option is set to allow administrative access. For more information about how to use these options, see the `na_options(1)` man pages.

Step

1. Point your browser to the following location:

```
http(s)://<system_name>/na_admin/logs/
```

`system_name` is the name of your storage system.

Related concepts

[The default security settings](#) on page 50

[How to access a storage system by using FilerView](#) on page 75

Related tasks

[Allowing only secure access to the storage system](#) on page 87

Changing the root volume

Every storage system must have a root volume. Therefore, you must always have one volume designated as the root volume. However, you can change which volume on your storage system is used as the root volume.

Before you begin

Before designating a volume to be the new root volume, ensure that the volume meets the minimum size requirement. 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, Data ONTAP prevents you from setting the root option.

Starting in Data ONTAP 8.0.1, you can use a 64-bit volume for the root volume.

Before designating a volume to be the new root volume, ensure that the volume has at least 2 GB of free space.

If you use a FlexVol volume for the root volume, ensure that it has a space guarantee of `volume`.

About this task

You might want to change the storage system's root volume, for example, when you migrate your root volume from a traditional volume to a FlexVol volume. 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 *Data ONTAP 7-Mode Storage Management Guide*.

Steps

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 *Data ONTAP 7-Mode Storage Management Guide*.

2. Using `ndmptcopy`, copy the `/etc` directory and all of its subdirectories from the current root volume to the new root volume. For more information about `ndmptcopy`, see the *Data ONTAP 7-Mode Data Protection Tape Backup and Recovery Guide*.
3. Enter the following command:

```
vol options vol_name root
```

`vol_name` is the name of the new root volume.

If the volume does not have at least 2 GB of free space, the command fails and an error message appears.

After a volume is designated to become the root volume, it cannot be brought offline or restricted.

Note: 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.

If you move the root volume outside the current root aggregate, you must also change the value of the aggregate `root` option (using `aggr options aggr_name root`) so that the aggregate containing the root volume becomes the root aggregate. Starting in Data ONTAP 8.0.1, you can use a 64-bit aggregate for the root aggregate. If you change the root aggregate, a new root volume is created during the subsequent boot only when the aggregate does not already contain a FlexVol designated as the root volume and when the aggregate has at least 2 GB of free space.

For more information about the aggregate `root` option, see the `na_aggr(1)` man page.

4. Enter the following command to reboot the storage system:

```
reboot
```

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.

Related concepts

Recommendations regarding the root volume on page 89

Size requirement for root FlexVol volumes on page 91

How to start and stop the storage system

You can start your storage system in several ways. You can boot the storage system from the system or boot environment prompt. You may also be able to start the storage system remotely. You can restart your system by halting and booting it.

Next topics

[How to boot the storage system](#) on page 103

[About rebooting the storage system](#) on page 110

[Halting the storage system](#) on page 112

How to boot the storage system

The storage system automatically boots Data ONTAP from a boot device, such as a PC CompactFlash Card. The system's boot device, shipped with the current Data ONTAP release and a diagnostic kernel, contains sufficient space for an upgrade kernel.

The storage system can be upgraded to the most recent Data ONTAP release. When you install new software, the `download` command copies a boot kernel to the boot device. For more information, see the *Data ONTAP 7-Mode Upgrade Guide*.

Next topics

[Ways to boot the storage system](#) on page 103

[Booting the storage system at the storage system prompt](#) on page 104

[Booting Data ONTAP at the boot environment prompt](#) on page 107

[Booting Data ONTAP remotely](#) on page 108

[Recovering from a corrupted image of the boot device](#) on page 109

[Checking available Data ONTAP versions](#) on page 110

Ways to boot the storage system

You can boot the storage system from the storage system prompt, the boot environment prompt, or the CLI prompt for the remote management device.

You can boot the storage system from the storage system prompt, for example, `toaster>`.

You can also boot Data ONTAP remotely from the CLI prompt of the remote management device, for example, `SP toaster>`, `RLM toaster>`, or `bmc shell ->`.

You can also boot the storage system with the following boot options from the boot environment prompt:

- `boot_ontap`

Boots the current Data ONTAP software release stored on the boot device (such as a PC 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 boot device 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 boot device. 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.
- `boot_diags`
Boots a Data ONTAP diagnostic kernel. For more information, see the *Diagnostics Guide*.

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

Note: Starting in Data ONTAP 8.0, netboot is not a supported function, unless you are restoring the Data ONTAP image on the boot device, such as a PC CompactFlash card. If you need to boot the storage system from a Data ONTAP image stored on a remote server, contact technical support. For information about how to replace a boot device or restore the Data ONTAP image on the boot device, see the *Replacing a boot device* flyer that is applicable to the version of Data ONTAP used by your platform.

Booting the storage system at the storage system prompt

The storage system is configured to boot from the boot device, such as a PC CompactFlash card. You can boot the storage system from the storage system prompt.

About this task

When you reboot the storage system, it reboots in normal mode by default. 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 initialize disks and reset system configuration for redeploying the storage system
- To restore configuration information back to the boot device

Steps

1. At the storage system prompt, enter the following command:

reboot

The storage system begins the boot process.

2. If you want the storage system to boot automatically in normal mode, allow the storage system to reboot uninterrupted.

The following message appears, indicating that you are done:

```
root logged in from console
```

3. If you want to select from a menu of alternative boot modes, press **Ctrl-C** to display the special boot menu when prompted to do so.

The storage system displays the following boot menu:

```
1) Normal Boot.
2) Boot without /etc/rc.
3) Change password.
4) Clean configuration and initialize all disks.
5) Maintenance mode boot.
6) Update flash from backup config.
7) Install new software first.
8) Reboot node.
Selection (1-8)?
```

4. Select one of the boot types by entering the corresponding number.

To ...	Select ...
Continue to boot the storage system normally	1) Normal Boot
Troubleshoot and repair configuration problems	2) Boot without /etc/rc. 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.
Change the password of the storage system	3) Change Password

To ...	Select ...
Initialize all the disks and create a FlexVol root volume	<p data-bbox="475 230 1150 253">4) <code>Clean configuration and initialize all disks</code></p> <p data-bbox="501 274 1204 328">Attention: This menu option erases all data on the disks and resets your system configuration to the factory default settings.</p> <p data-bbox="501 348 1228 487">If you need to preserve existing configuration values (such as your system IP address, gateway addresses, and DNS server addresses) that are used for system setup, make a note of the values before selecting this menu option. You can find your current setup settings by entering <code>setup</code> at the storage system prompt.</p> <p data-bbox="475 526 1177 548">This menu option reboots the storage system before initializing the disks.</p> <p data-bbox="475 569 1237 678">For a V-Series system that has a disk shelf, this menu option initializes only the disks on the disk shelf, not the array LUNs. For a V-Series system that does not have a disk shelf, this menu option initializes the root volume on the storage array.</p> <p data-bbox="475 699 1208 779">After the initialization procedure has finished, the setup script starts and prompts you for configuration information. For information about setting up the storage system, see the <i>Data ONTAP 7-Mode Software Setup Guide</i>.</p> <p data-bbox="501 800 1150 878">Note: Data ONTAP 8.0 or later does not allow you to create a new traditional root volume from the boot menu. However, preexisting traditional root volumes are still supported.</p>
Perform some aggregate and disk operations and get detailed aggregate and disk information.	<p data-bbox="475 913 817 935">5) <code>Maintenance mode boot</code></p> <p data-bbox="501 956 1099 979">Note: Maintenance mode is special for the following reasons:</p> <ul data-bbox="501 999 1237 1152" style="list-style-type: none"> <li data-bbox="501 999 1208 1022">• Most normal functions, including file system operations, are disabled. <li data-bbox="501 1034 1237 1088">• A limited set of commands is available for diagnosing and repairing disk and aggregate or volume problems. <li data-bbox="501 1100 1177 1152">• You exit Maintenance mode with the <code>halt</code> command. To reboot the storage system, enter <code>boot</code> after the firmware prompt.
Restore the configuration information from the root volume to the boot device, such as a PC CompactFlash card	<p data-bbox="475 1204 962 1227">6) <code>Update flash from backup config</code></p> <p data-bbox="501 1248 1228 1413">Note: Data ONTAP stores some system configuration information on the boot device. When the storage system boots, the information on the boot device is automatically backed up onto the root volume. If the boot device becomes corrupted or needs to be replaced, you use this menu option to restore the configuration information from the root volume back to the boot device.</p>

To ...	Select ...
Install new software on a V-Series system	<p>7) Install new software first</p> <p>If the Data ONTAP software on the boot device does not include support for the storage array you want to use for the root volume, you can use this menu option to obtain a version of the software that supports your storage array and install it on your system.</p> <p>Note: This menu option is only for installing a newer version of Data ONTAP software on a V-Series system that has no root volume installed. Do <i>not</i> use this menu option to upgrade the Data ONTAP software on either a FAS system or a V-Series system.</p>
Reboot the storage system	8) Reboot node

For additional information about the boot menu, see the `na_floppyboot(1)` man page.

Booting Data ONTAP at the boot environment prompt

You can boot the current release or the backup release of Data ONTAP when you are at the boot environment prompt.

Steps

1. If you are at the storage system prompt, enter the following command:

```
halt
```

The storage system console displays the boot environment prompt.

2. At the boot environment prompt, enter one of the following commands:

To boot...	Enter...
The current release of Data ONTAP	<code>boot_ontap</code>
The Data ONTAP primary image from the boot device	<code>boot_primary</code>
The Data ONTAP backup image from the boot device	<code>boot_backup</code>

Note: For more information about commands available from the boot prompt, enter `help` at the firmware prompt for a list of commands or `help command` for command details.

Booting Data ONTAP remotely

You can boot Data ONTAP remotely by using the remote management device.

Steps

1. From the administration host, log in to the remote management device by entering the following command:

```
ssh username@IP_for_remote_management_device
```

The CLI prompt for the remote management device, which can be one of the following depending on the storage system model, appears.

```
SP toaster>
RLM toaster>
bmc shell>
```

2. To turn on the storage system, enter the following command at the CLI prompt for the remote management device:

```
system power on
```

3. To access the system console, enter the following command at the CLI prompt for the remote management device:

```
system console
```

The storage system prompt appears.

4. If the storage system does not boot automatically, enter one of the following commands at the boot environment prompt:

To boot...	Enter...
The current release of Data ONTAP	boot_ontap
The Data ONTAP primary image from the boot device	boot_primary
The Data ONTAP backup image from the boot device	boot_backup

Related concepts

[Ways to boot the storage system](#) on page 103

[Managing storage systems remotely](#) on page 189

[Prerequisites for logging in to the SP](#) on page 0

[How to log in to the RLM](#) on page 221

[How to log in to the BMC](#) on page 253

Recovering from a corrupted image of the boot device

You can recover from a corrupted image of the boot device (such as the CompactFlash card) for a storage system by using the remote management device.

Steps

1. Log in to the remote management device by entering the following command at the administration host:

```
ssh username@IP_for_remote_management_device
```

The CLI prompt for the remote management device, which can be one of the following depending on the storage system model, appears.

```
SP toaster>
RLM toaster>
bmc shell>
```

2. At the CLI prompt for the remote management device, perform one of the following steps:
 - To reboot the storage system by using the primary image, enter the following command:

```
system reset primary
```

- To reboot the storage system by using the backup image, enter the following command:

```
system reset backup
```

Note: For the BMC, the `system reset` command is an advanced command. You should use advanced commands only when technical support directs you to do so.

The following prompt is displayed:

```
This will cause a dirty shutdown of your appliance. Continue? [y/n]
```

3. Enter `y` to continue.

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.

Related concepts

[Ways to boot the storage system](#) on page 103

[Managing storage systems remotely](#) on page 189

[Prerequisites for logging in to the SP](#) on page 0

[How to log in to the RLM](#) on page 221

[How to log in to the BMC](#) on page 253

Checking available Data ONTAP versions

You might need to check the current booted kernel and other kernels available on the boot device (such as the CompactFlash card) if an upgrade was unsuccessful or if you need to run kernel diagnostics.

About this task

By default, the storage system boots the current Data ONTAP release from the primary kernel.

Step

1. Do one of the following:

To determine...	At the storage system console, enter...
The current booted Data ONTAP version	<code>version</code>
Data ONTAP versions available on the boot device	<code>version -b</code>

If you enter `version`, the console displays the version number of Data ONTAP that is currently running.

If you enter `version -b`, the console displays information from the boot device, 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.

About rebooting the storage system

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.

Next topics

[Rebooting the storage system from the system console](#) on page 111

[Rebooting the storage system remotely](#) on page 111

Rebooting the storage system from the system console

You can reboot the storage system if the system console is displaying the command prompt.

Steps

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:

```
reboot [-t minutes]
```

-t minutes is the amount of time that elapses before the reboot occurs.

Rebooting the storage system remotely

You can reboot your storage system remotely by using the remote management device.

Steps

1. From the administration host, log in to the remote management device by entering the following command:

```
ssh username@IP_for_remote_management_device
```

The CLI prompt for the remote management device, which can be one of the following depending on the storage system model, appears.

```
SP toaster>
RLM toaster>
bmc shell>
```

2. At the CLI prompt for the remote management device, enter the following command to access the system console:

```
system console
```

The storage system prompt appears.

```
toaster>
```

3. At the storage system prompt, enter the following command to reboot the storage system:

```
reboot
```

Related concepts

[Ways to boot the storage system](#) on page 103

[Managing storage systems remotely](#) on page 189

Prerequisites for logging in to the SP on page 0

How to log in to the RLM on page 221

How to log in to the BMC on page 253

Halting the storage system

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

About this task

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 (using `cifs terminate`) without giving CIFS users enough time to save their changes can cause data loss.

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.

Step

1. Enter the following command:

```
halt [-d dump_string] [-t interval] [-f]
```

`-d dump_string` causes the storage system to perform a core dump before halting. You use `dump_string` to describe the reason for the core dump. The message for the core dump will include the reason specified by `dump_string`.

Attention: Using `halt -d` causes an improper shutdown of the storage system (also called a dirty shutdown). Avoid using `halt -d` for normal maintenance shutdowns. For more details, see the `na_halt(1)` man page.

`-t interval` causes the storage system to halt after the number of minutes you specify for the interval.

`-f` prevents one partner in a high-availability configuration from taking over the other after the storage system halts.

The storage system displays the boot prompt. When you see the boot prompt, you can turn the power off.

How to manage administrator and diagnostic access

Data ONTAP enables you to control access to your storage system to provide increased security and auditing capability. It also enables you to manage passwords on the storage system to ensure security.

Next topics

[Reasons for creating administrator accounts](#) on page 115

[Root access to the storage system](#) on page 118

[How to manage users](#) on page 120

[How to manage groups](#) on page 123

[How to manage roles](#) on page 127

[Users, groups, and roles](#) on page 133

[Administrative user creation examples](#) on page 138

[How to manage passwords for security](#) on page 140

[The diagnostic account and the systemshell](#) on page 146

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.

The following are the reasons for creating administrator accounts:

- 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 audit-log 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 storage system using Telnet, RSH, SSH, FilerView, Data ONTAP APIs, and Windows Remote Procedure Calls (RPCs).

For more information about authenticating users using Windows domains, see the section on user accounts in the CIFS chapter of the *Data ONTAP 7-Mode File Access and Protocols Management Guide*.

Next topics

[What users, groups, roles, and capabilities are](#) on page 116

[How users are assigned capabilities](#) on page 117

[Requirements for naming users, groups, and roles](#) on page 117

[Windows special groups](#) on page 117

[About changing capabilities of other groups and roles](#) on page 118

What users, groups, roles, and capabilities are

You need to understand what users, groups, roles, and capabilities are, so that you can grant different levels of administrative access to users of a storage system.

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 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.

Requirements for naming users, groups, and roles

When you name your users, groups and roles, you must meet the naming requirements.

The naming requirements are as follows:

- Names are case insensitive.
- Names can contain any alphanumeric character, a space, or a symbol that is not one of the following characters:
" * + , / \ : ; < = > ? []

Note: If the name contains spaces or special characters, enclose the name in double quotes (" ") when you use it in a command.

- You cannot give a user and a group the same name.

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

About 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 changes to that other group.

The changes you can make include the following:

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

Root access to the storage system

By default, root access to the storage system is enabled. You have the option to disable the root account's access to the storage system, preventing the root account from logging in the system or executing any commands.

To prevent the root account from logging in to the system or executing any commands, a user other than root with the `security-complete-user-control` security capability can disable root access by setting the option `security.passwd.rootaccess.enable` to `off`.

An EMS message is sent every time the option changes.

To reset the `security.passwd.rootaccess.enable` option to `on` (the default) to enable root access, a user must change the root account's password.

The option to enable or disable root access is supported if you access the storage system through Telnet, RSH, SSH, http-admin, NDMP, or the serial console.

Next topics

[Disabling root access to the storage system](#) on page 119

[Displaying the status of root access](#) on page 119

Related concepts

[Supported capability types](#) on page 129

Disabling root access to the storage system

Disabling the root account's access to the storage system prevents the root account from logging in to the system or executing any commands.

About this task

You can disable root access if you are a user other than root with the `security-complete-user-control` security capability, and if you access the storage system through Telnet, RSH, SSH, HTTP Admin, NDMP, or the serial console.

The storage system's root account is mapped to the `naroot` account of the remote management device (which can be the SP, the RLM, or the BMC.) If you disable the root account's access to the storage system, the `naroot` access to the storage system is automatically disabled if your system uses the SP or the RLM. However, disabling root access is not supported if your system uses the BMC.

Step

1. Enter the following command:

```
options security.passwd.rootaccess.enable off
```

The default is on.

Note: To reset the `security.passwd.rootaccess.enable` option to on to enable root access, a user other than root must first change the root account's password.

Displaying the status of root access

The status of the root account shows whether its access to the storage system is currently enabled.

Step

1. Enter one of the following commands:

```
options security.passwd.rootaccess.enable
useradmin user list root
```

Examples of root access status display

The following examples show that root access is currently disabled.

```
toaster> options security.passwd.rootaccess.enable
security.passwd.rootaccess.enable off
```

```
toaster> useradmin user list root
Name: root
Info: Default system administrator.
Rid: 0
Groups:
```

```
Full Name:
Allowed Capabilities: *
Password min/max age in days: 0/never
Status: disabled
```

How to manage users

You can create users, grant them access to the storage system, and modify their capabilities.

Next topics

[Creating users and assigning them to groups](#) on page 120

[Granting access to Windows domainusers](#) on page 122

[How to grant permissions for MMC](#) on page 123

[About changing another user's capabilities](#) on page 123

Creating users and assigning them to groups

You can create or modify a user and assign that user to one or more predefined or customized groups, giving that user the roles and capabilities associated with those groups.

About this task

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.”

You can have a maximum of 96 administrative users on a storage system.

Steps

1. Enter the following command:

```
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]
```

- Use `useradmin user add` to create a new user. Use `useradmin user modify` to modify the attributes of an existing user.
- `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. If `user_name` contains a space, enclose `user_name` in double quotes (" ").
- `comments` specifies a maximum 128-character comment that can be viewed through the `useradmin user 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.

To assign a user to the Compliance Administrators group, ensure that the `telnet.distinct.enable` option is set to `on`.

- *password_min_age* specifies the minimum number of days that users must have a password before they can change it. The default value is 0. If you specify a value larger than 4,294,967,295, the value is set to 4,294,967,295.
- *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. If you specify a value larger than 4,294,967,295, the value is set to 4,294,967,295. The password expires at midnight in the GMT time zone, on the expiration date.

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 of user creation

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
```

Related concepts

[Predefined groups](#) on page 124

[Requirements for naming users, groups, and roles](#) on page 117

Related tasks

[Assigning roles to groups by creating or modifying a group](#) on page 125

Granting access to Windows domainusers

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.

About this task

By default, the domain administrator account has full access to the system. You can log in this account by using the *domain\administrator* format with the appropriate password.

Steps

1. To assign a Windows domainuser to a custom or predefined group, enter the following command:

```
useradmin domainuser add win_user_name -g {custom_group|
Administrators|"Backup Operators"|Guests|"Power Users"|Users}[,...]
```

win_user_name is the Windows domainuser 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.

Administrators | "Backup Operators" | Guests | "Power Users" | Users are groups predefined by Data ONTAP with default roles and capabilities.

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"
```

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.

Related concepts

[How to manage users](#) on page 120

[Predefined groups](#) on page 124

How to grant permissions for MMC

In order to use Microsoft Management Console (MMC) to access the storage system, a user must be in the local Administrators group. Because the Domain Admins group is placed within the Administrators group, users in the Domain Admins group have MMC access also.

The following are the methods for adding users to the Administrators group for MMC access:

- Add local users (users that were created on the storage system) by using the `useradmin user modify username -g Administrators` command.
- Add nonlocal users (users that exist on the domain) by using the `useradmin domainuser add domain\username -g Administrators` command.
- Use the MMC on the domain to add `domain\username` to the Domain Admins group.

Related tasks

[Creating users and assigning them to groups](#) on page 120

[Granting access to Windows domainusers](#) on page 122

About 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 change another user's capabilities or account information.

The changes you can make include:

- 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.

How to manage groups

You can use groups predefined by Data ONTAP or create or modify a group.

Next topics

Predefined groups on page 124

Assigning roles to groups by creating or modifying a group on page 125

Renaming a group on page 126

Loading groups from the lclgroups.cfg file on page 126

Setting the maximum number of auxiliary UNIX groups allowed for a user on page 127

Predefined groups

You can assign a user or domainuser to a predefined set of groups and roles provided by Data ONTAP. The predefined groups include Administrators, Power Users, Compliance Administrators, Backup Operators, Users, Guests, and Everyone.

The following table describes the predefined groups.

Predefined Group	Default roles	Default privileges
Administrators	admin	Grants all CLI, API, login, and security capabilities.
Power Users	power	Grants the ability to perform the following tasks: <ul style="list-style-type: none"> • Invoke all <code>cifs</code>, <code>exportfs</code>, <code>nfs</code>, and <code>useradmin</code> CLI commands • Make all <code>cifs</code> and <code>nfs</code> API calls • Log in to Telnet, HTTP, RSH, and SSH sessions
Compliance Administrators	compliance	Grants the ability to execute compliance-related operations. <p>Note: You cannot assign a user to this group if the <code>telnet.distinct.enable</code> option is set to <code>off</code>.</p>
Backup Operators	backup	Grants the ability to make NDMP requests.
Users	audit	Grants the ability to make <code>snmp-get</code> and <code>snmp-get-next</code> API calls.
Guests	none	None
Everyone	none	None

Related concepts

[Predefined roles](#) on page 128

[Supported capability types](#) on page 129

Assigning roles to groups by creating or modifying a group

You can create or modify a group, giving that group the capabilities associated with one or more predefined or customized roles.

About this task

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.

Steps

1. Use the `useradmin group add` command to create a new group or the `useradmin group modify` command to modify a group, by entering the following command:

```
useradmin group {add|modify} group_name [-c comments] [-r {custom_role|
root|admin|power|backup|compliance|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.

Note: Do not create groups with the same name as any of the Windows special groups or any existing users.

custom_role is a customized role with capabilities assigned through the `useradmin role add` command.

`root`, `admin`, `power`, `backup`, `compliance`, and `audit` are roles predefined with default capabilities by Data ONTAP.

Example

The following command gives the group “admin users” capabilities associated with the `admin` role, and removes any roles previously assigned to the `admin_users` group.

```
useradmin group modify "admin users" -r admin
```

2. Enter the following command to verify the success of your operation:

```
useradmin group list group_name
```

The roles and capabilities assigned to the group in question are listed in the output of this command.

Related concepts

[Requirements for naming users, groups, and roles](#) on page 117

[Windows special groups](#) on page 117

[Predefined roles](#) on page 128

Renaming a group

You can change the name of an existing group.

Step

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.

Related concepts

[Windows special groups](#) on page 117

Loading groups from the lclgroups.cfg file

When groups are created, they are placed in the `lclgroups.cfg` file. Normally, this file is for administrative reference only. It is not used to reload groups into the system memory. However, sometimes you need Data ONTAP to reload this file, for example, when you are migrating a storage system or a vFiler unit.

About this task

Using this procedure unloads the current groups from memory before loading the new file; currently configured groups will no longer be available unless they are also configured in the new file.

To perform this operation, the user must belong to a group that has the `security-load-lclgroups` capability.

Do not edit the `lclgroups.cfg` file directly to add or remove groups. Use the `useradmin group` command to administer groups.

Steps

1. Using a client, copy the new `lclgroups.cfg` file to the `/etc` directory, giving it a different name.
2. Enter the following command:

```
useradmin domainuser load new_lclgroups.cfg_filename
```

new_lclgroups.cfg_filename is the name of the new `lclgroups.cfg` file you created in Step 1.

The groups in the current `lclgroups.cfg` file are unloaded from memory and the groups in the new `lclgroups.cfg` file are loaded into memory. In addition, the current `lclgroups.cfg` file is moved to `lclgroups.cfg.bak`, and a new `lclgroups.cfg` file is created from the file you specified.

Setting the maximum number of auxiliary UNIX groups allowed for a user

If you use Kerberos V5 authentication, the maximum number of auxiliary UNIX groups that a user can be a member of is 32 by default. You can increase the maximum to 256 groups by setting the `nfs.max_num_aux_groups` option to 256.

About this task

If you do not use Kerberos V5 authentication, the maximum number of auxiliary UNIX groups that a user can be a member of is 16.

Step

1. To change the maximum number of auxiliary UNIX groups that a user can be a member of, enter the following command:

```
options nfs.max_num_aux_groups [32 | 256]
```

The default value is 32.

Note: In FlexCache setups, Data ONTAP supports a maximum of 32 auxiliary UNIX groups for FlexCache volumes, regardless of the value set for this option.

For more information about the `nfs.max_num_aux_groups` option, see the `na_options(1)` man page.

How to manage roles

You can use roles predefined by Data ONTAP or create new roles. You can also modify an existing role.

Next topics

[Predefined roles](#) on page 128

[Supported capability types](#) on page 129

[Creating a new role and assigning capabilities to roles](#) on page 131

[Modifying an existing role or its capabilities](#) on page 132

Predefined roles

The predefined roles Data ONTAP provides include `root`, `admin`, `power`, `backup`, `compliance`, `audit`, and `none`.

The following table describes the roles that are predefined by Data ONTAP.

Role	Default capability assignments	Summary of default granted capabilities
<code>root</code>	<code>*</code>	Grants all possible capabilities.
<code>admin</code>	<code>cli-*</code> , <code>api-*</code> , <code>login-*</code> , <code>security-*</code>	Grants all CLI, API, login, and security capabilities.
<code>power</code>	<code>cli-cifs*</code> , <code>cli-exportfs*</code> , <code>cli-nfs*</code> , <code>cli-useradmin*</code> , <code>api-cifs-*</code> , <code>api-nfs-*</code> , <code>login-telnet</code> , <code>login-http-admin</code> , <code>login-rsh</code> , <code>login-ssh</code> , <code>api-system-api-*</code>	Grants the ability to : <ul style="list-style-type: none"> • Invoke all <code>cifs</code>, <code>exportfs</code>, <code>nfs</code>, and <code>useradmin</code> CLI commands • Make all <code>cifs</code> and <code>nfs</code> API calls • Log in using Telnet, HTTP, RSH, and SSH sessions
<code>backup</code>	<code>login-ndmp</code>	Grants the ability to make NDMP requests.
<code>compliance</code>	<code>cli-cifs*</code> , <code>cli-exportfs*</code> , <code>cli-nfs*</code> , <code>cli-useradmin*</code> , <code>api-cifs-*</code> , <code>api-nfs-*</code> , <code>login-telnet</code> , <code>login-http-admin</code> , <code>login-rsh</code> , <code>login-ssh</code> , <code>api-system-api-*</code> , <code>cli-snaplock*</code> , <code>api-snaplock-*</code> , <code>api-file-*</code> , <code>compliance-*</code>	Grants compliance-related capabilities in addition to all the capabilities granted by the <code>power</code> role. <p>Note: The <code>compliance</code> role is the default role for the Compliance Administrators group. The <code>compliance</code> role cannot be removed from the Compliance Administrators group or added to other groups.</p>
<code>audit</code>	<code>api-snmp-get</code> , <code>api-snmp-get-next</code>	Grants the ability to make <code>snmp-get</code> and <code>snmp-get-next</code> API calls.
<code>none</code>	None	Grants no administrative capabilities.

Related concepts

[Predefined groups](#) on page 124

[Supported capability types](#) on page 129

Related tasks

[Assigning roles to groups by creating or modifying a group](#) on page 125

Supported capability types

The capability types Data ONTAP supports include `login`, `cli`, `security`, `api`, `compliance`, and `filerview`.

The following table describes the supported capability types.

This capability type...	Has the following capabilities...
login	<p>Grants the specified role login capabilities.</p> <p><code>login-*</code> grants the specified role the capability to log in through all supported protocols.</p> <p><code>login-protocol</code> grants the specified role the capability to log in through a specified protocol. Supported protocols include the following:</p> <ul style="list-style-type: none"> • <code>login-console</code> grants the specified role the capability to log in to the storage system using the console. • <code>login-http-admin</code> grants the specified role the capability to log in to the storage system using HTTP. • <code>login-ndmp</code> grants the specified role the capability to make NDMP requests. • <code>login-rsh</code> grants the specified role the capability to log in to the storage system using RSH. • <code>login-snmp</code> grants the specified role the capability to log in to the storage system using SNMPv3. • <code>login-sp</code> grants the specified role the capability to log in to the SP or the RLM by using SSH. • <code>login-ssh</code> grants the specified role the capability to log in to the storage system using SSH. • <code>login-telnet</code> grants the specified role the capability to log in to the storage system using Telnet.
cli	<p>Grants the specified role the capability to execute one or more Data ONTAP command line interface (CLI) commands.</p> <p><code>cli-*</code> grants the specified role the capability to execute all supported CLI commands.</p> <p><code>cli-cmd*</code> grants the specified role the capability to execute all commands associated with the CLI command <code>cmd</code>.</p> <p>For example, the following command grants the specified role the capability to execute all <code>vol</code> commands:</p> <pre>useradmin role modify status_gatherer -a cli-vol*</pre> <p>Note: Users with <code>cli</code> capability also require at least one <code>login</code> capability to execute CLI commands.</p>

This capability type...	Has the following capabilities...
<p>security</p>	<p>Grants the specified role security-related capabilities, such as the capability to change other users' passwords or to invoke the CLI <code>priv set advanced</code> command.</p> <p><code>security-*</code> grants the specified role all security capabilities.</p> <p><code>security-capability</code> grants the specified role one of the following specific security capabilities:</p> <ul style="list-style-type: none"> • <code>security-api-vfiler</code> grants the specified role the capability to forward or tunnel ONTAP APIs from the physical storage system into a vFiler unit for execution. • <code>security-passwd-change-others</code> grants the specified role the capability to change the passwords of all users with equal or fewer capabilities. • <code>security-priv-advanced</code> grants the specified role the capability to access the advanced CLI commands. • <code>security-load-lclgroups</code> grants the specified role the capability to reload the <code>lclgroups.cfg</code> file. • <code>security-complete-user-control</code> grants the specified role the capability to create, modify, and delete users, groups, and roles with greater capabilities.
<p>api</p>	<p>Grants the specified role the capability to execute Data ONTAP API calls.</p> <p><code>api-*</code> grants the specified role all API capabilities.</p> <p><code>api-api_call_family-*</code> grants the specified role the capability to call all API routines in the family <code>api_call_family</code>.</p> <p><code>api-api_call</code> grants the specified role the capability to call the API routine <code>api_call</code>.</p> <p>Note:</p> <p>You have more fine-grained control of the command set with the <code>api</code> capabilities because you can give subcommand capabilities as well.</p> <p>Users with <code>api</code> capability also require the <code>login-http-admin</code> capability to execute API calls.</p>
<p>compliance</p>	<p>Grants the specified role the capability to execute compliance-related operations.</p> <p><code>compliance-*</code> grants the specified role the capability to execute all compliance-related operations.</p> <p><code>compliance-privileged-delete</code> grants the specified role the capability to execute privileged deletion of compliance data.</p> <p>Note: The compliance capabilities (<code>compliance-*</code>) are included in the default capabilities of the <code>compliance</code> role. The compliance capabilities cannot be removed from the <code>compliance</code> role or added to other roles.</p>

This capability type...	Has the following capabilities...
filerview	<p>Grants the specified role read-only access to FilerView.</p> <p>This capability type includes only the <code>filerview-readonly</code> capability, which grants the specified role the capability to view but not change manageable objects on systems managed by FilerView.</p> <p>Note:</p> <p>There is no predefined role or group for read-only FilerView access. You must first assign the <code>filerview-readonly</code> capability to a role and then assign the role to a group, before you can create a user in such a group.</p>

Related concepts

[About changing another user's capabilities](#) on page 123

[Predefined roles](#) on page 128

[Predefined groups](#) on page 124

Related tasks

[Loading groups from the `lclgroups.cfg` file](#) on page 126

[Creating a new role and assigning capabilities to roles](#) on page 131

[Assigning roles to groups by creating or modifying a group](#) on page 125

Creating a new role and assigning capabilities to roles

You can create a new role and grant desired capabilities to the role.

Steps

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.

comments is a short string you can use to document this role.

The *capability* parameters are the types of access you want to grant to this new role.

Example

You can also 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-*
```

2. To verify the success of the operation, enter the following command:

```
useradmin role list role_name
```

The capabilities allowed for the specified role are listed.

Related concepts

[About changing another user's capabilities](#) on page 123

[Requirements for naming users, groups, and roles](#) on page 117

Modifying an existing role or its capabilities

You can modify an existing role's capabilities or its comments.

About this task

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.

Steps

1. Enter the following command:

```
useradmin role modify role_name [-c comments] -a  
capability1[,capability2...] [-f]
```

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 types of access you want to grant to this role.

The `-f` option forces the change without a warning.

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
```

The capabilities allowed for the specified role are listed.

Users, groups, and roles

You can display information for existing users, groups, or roles. You can also delete them.

Next topics

[Commands that list users, domainusers, groups, or roles](#) on page 133

[Commands that delete users, domainusers, groups, or roles](#) on page 137

Commands that list users, domainusers, groups, or roles

You use the `useradmin` commands to display information for users, domainusers, groups, or roles.

The following table describes the commands.

Command	Description
<code>useradmin whoami</code>	Displays the user name of the account you are currently using.
<code>useradmin user list</code>	Lists all administrative users configured for this storage system. Each user entry includes the user name, comment information, a user ID number generated by Data ONTAP, and groups that each user belongs to.
<code>useradmin user list user_name</code>	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 user ID number generated by Data ONTAP, effective allowed capabilities, and user account status.
<code>useradmin user list -x</code>	Lists the extended information for all administrators. 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 user ID number generated by Data ONTAP, effective allowed capabilities, and user account status.
<code>useradmin user list -g grp_name</code>	Lists information for all users assigned to a specified group.

Command	Description
<code>useradmin domainuser list -g <i>group_name</i></code>	<p>Lists the SIDs of all Windows domain administrative users assigned to a specified group.</p> <p>To list the user name, comment information, and the groups that each user belongs to, follow up with <code>cifs lookup</code> and <code>useradmin user list</code> commands.</p> <p>Note: The Rid value of 500 for the Administrator user corresponds to the last number in the Administrator user's SID.</p>
<code>useradmin group list</code>	Lists all the administrative user groups configured for this storage system. Each group entry includes the group name, comment information, user ID number generated by Data ONTAP, and every role associated with that group.
<code>useradmin group list <i>group_name</i></code>	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.
<code>useradmin role list</code>	Lists all the roles configured for this storage system. Each role entry lists the role name, comment information, and allowed capabilities.
<code>useradmin role list <i>role_name</i></code>	Lists the information for a single specified role name.

Example of `useradmin whoami` output

```
toaster> useradmin whoami
Administrator
```

Example of `useradmin user list` output

```
toaster> 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
...
```

Example of `useradmin user list user_name` output

```
toaster> 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
```

Example of `useradmin user list -x` output

```
toaster> useradmin user list -x
Name: administrator
Info: Built-in account for administering the filer
Rid: 500
Groups: Administrators
Full Name:
Allowed Capabilities: login-*,cli-*,api-*,security-*
Password min/max age in days: 0/4294967295
Status: enabled

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
...
```

Example of `useradmin user list -g grp_name` output

```
toaster> useradmin user list -g Administrators
Name: administrator
Info: Built-in account for administering the filer
Rid: 500
Groups: Administrators

Name: marshall
Info:
Rid: 131454
Groups: Administrators
```

```
...
```

Example of `useradmin domainuser list -g group_name` output

```
toaster> 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.

toaster> cifs lookup S-1-7-24-1214340929-620487827-8395249115-512
name = MBS-LAB\Domain Admins

toaster> cifs lookup S-1-7-24-1838915891-154599588-1081798244-500
name = ZND\Administrator

toaster> 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-*
```

Example of `useradmin group list` output

```
toaster> 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
...
```

Example of `useradmin group list group_name` output

```
toaster> useradmin group list Administrators
Name: Administrators
Info: Members can fully administer the filer.
Rid: 544
```

```
Roles: admin
Allowed Capabilities: login-*,cli-*,api-*,security-*
```

Example of useradmin role list output

```
toaster> 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:

...
```

Example of useradmin role list role_name output

```
toaster> useradmin role list admin
Name:      admin
Info:      Default role for administrator privileges.
Allowed Capabilities: login-*,cli-*,api-*,security-*
```

Commands that delete users, domainusers, groups, or roles

You use the useradmin commands to delete users, domainusers, groups, or roles.

The following table describes the commands.

Command	Description
<code>useradmin user delete user_name</code>	<p>Deletes the specified user from the storage system.</p> <p>The <code>useradmin user delete</code> command deletes any local user except for root. User names are case insensitive.</p> <p>Note: You cannot delete or modify a user with greater capabilities than you have.</p>

Command	Description
<code>useradmin domainuser delete win_user_name -g group1,[group2,...]</code>	Removes the specified user from the specified group or groups. User names are case insensitive. 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 <i>useradmin user delete</i> command instead.
<code>useradmin group delete group_name</code>	Deletes the specified group from the storage system. Group names are case insensitive. Note: All users must be removed from a group before the group itself can be deleted.
<code>useradmin role delete role_name</code>	Deletes the specified role from the storage system. Role names are case insensitive. Note: A role that is still assigned to a group cannot be deleted.

Administrative user creation examples

You can create a user with custom capabilities or no administrative capabilities, thereby controlling the user's administrative access.

Next topics

[Example of creating a user with custom capabilities](#) on page 138

[Example of creating a user with no administrative capabilities](#) on page 139

Example of creating a user with custom capabilities

You can 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`.

```
toaster> useradmin role add only_ssh -a login-ssh
Role <only_ssh> added.
Thu Apr 22 10:50:05 PDT [toaster: useradmin.added.deleted:info]: The role
'only_ssh' has been added.

toaster> useradmin role add qtree_commands -a cli-qtree*,api-qtree-*
Role <qtree_commands> added.
Thu Apr 22 10:51:51 PDT [toaster: useradmin.added.deleted:info]: The role
'qtree_commands' has been added.

toaster> useradmin group add ssh_qtree_admins -r only_ssh,qtree_commands
Group <rsh_qtree_admins> added.
Thu Apr 22 10:53:07 PDT [toaster: useradmin.added.deleted:info]: The group
'ssh_qtree_admins' has been added.

toaster> useradmin user add wilma -g ssh_qtree_admins
New password:
Retype new password:
User <wilma> added.
Thu Apr 22 10:54:43 PDT [toaster: useradmin.added.deleted:info]: The user
'wilma' has been added.

toaster> useradmin user list wilma
Name: wilma
Info:
Rid: 131074
Groups: ssh_qtree_admins
Full Name:
Allowed Capabilities: login-ssh,cli-qtree*,api-qtree-*
```

Example of creating 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.

Steps

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
```

“Allowed Capabilities” should be blank.

How to manage passwords for security

Data ONTAP provides several methods you can use to ensure that the password policies for your storage systems meet your company's security requirements.

The following are the methods you can use:

- Password rules

The `security.passwd.rules` options enable you to specify rules for valid passwords.

The following are the default password rules for all accounts when `security.passwd.rules.enable` is set to on (the default):

- The password must be at least eight characters long.
- The password must contain at least one number.
- The password must contain at least two alphabetic characters.

Note: During the initial setup of a storage system shipped with Data ONTAP 8.0 or later, you are prompted to set up a password for the root account by following these password rules. Subsequent invocations of the `setup` command do not prompt you to set up a password for the root account.

For more information about setting up the storage system, see the *Data ONTAP 7-Mode Software Setup Guide*.

You can change the password rules by using the `security.passwd.rules` options. For more information about the password rule options, see the `na_options(1)` man page.

- Password history

The password history feature enables you to require users to create new passwords that are different from a specified number of previously used passwords, rather than simply using the same password every time. You use the `security.passwd.rules.history` option to specify how many unique passwords users need to create before they can reuse a password.

For storage systems shipped with Data ONTAP 8.0 or later, the default value is 6. In this case, the password a user creates cannot be the same as any of that user's last six passwords.

For storage systems upgraded to Data ONTAP 8.0 or later from an earlier release, the setting for the `security.passwd.rules.history` option stays the same as before the upgrade.

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

- Password expiration (maximum age)

The password expiration feature enables you to require that users change their passwords before they have had the password for the specified number of days. You use the `-M` option of the `useradmin user add` or the `useradmin user modify` command to specify the maximum password duration for individual users. The default value is 4,294,967,295. For more information, see the `na_useradmin(1)` man page.

Note: Before using the password expiration feature, 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.

- Password minimum age

Password minimum age (a specified minimum length of time each password stays in effect) prevents users from changing their passwords too soon, thus cycling through their previous passwords too quickly. You use the `-m` option of the `useradmin user add` or the `useradmin user modify` command to specify the minimum password duration for individual users. The default value is 0, which does not enforce a minimum password age. For more information, see the `na_useradmin(1)` man page.

Note: Before using the password minimum age feature, 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.

- Password lockout

The password lockout feature enables you to lock out users (except the root account) 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 the number of tries a user can make before being locked out of the system. The default value is 4,294,967,295. 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 an administrator 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.

Next topics

[Changing the storage system password](#) on page 141

[Changing a local user account password](#) on page 142

[Options that manage password rules](#) on page 143

Changing the storage system password

You can change the storage system password, which is also the password for the root user account.

About this task

The `naroot` user account, which can be used to log in to the remote management device, uses the storage system root password. Changing the storage system password also changes the password for `naroot`.

Step

1. Do one of the following:

If you are using this connection method to administer the storage system...	Then...
Telnet session or the console	<ol style="list-style-type: none"> Enter the following command at the storage system prompt: <code>passwd</code> Enter the storage system account name: <code>root</code> Enter the existing storage system password (not required if you are root or have the <code>security-passwd-change-others</code> capability). Enter a new password, and then enter it a second time to confirm it.
Remote Shell connection	Enter the following command from a UNIX host: <pre>rsh system_name -l root:root_password passwd old_password new_password root</pre>
Secure Shell connection	Enter the following command from a UNIX host: <pre>ssh -l root system_name passwd old_password new_password root</pre>

Related concepts

[The default security settings](#) on page 50

Changing a local user account password

You can change a local user account password by using a Telnet session, the console, the Secure Shell connection, or the Remote Shell connection.

Step

1. Do one of the following :

If you are using this connection method to administer the storage system... **Then...**

Telnet session or the console

- a. Enter the following command:
`passwd`
- b. When Data ONTAP prompts you, enter the name of the local user whose password you want to change.
- c. When Data ONTAP prompts you, enter the new password.
- d. Enter the new password again for confirmation.

Remote Shell connection

Enter the following command:

```
rsh system_name -l username:password passwd
old_password new_password username
```

Secure Shell connection

Enter the following command:

```
ssh -l username system_name passwd old_password
new_password username
```

Related concepts

[The default security settings](#) on page 50

Options that manage password rules

Data ONTAP provides the options to control password rules. Using the `options` command, you can specify password requirements such as how a check for password composition is performed and what the maximum or minimum number of characters is for a password.

The following table describes the options you can use to manage password rules.

Password rule option (used with the <code>options</code> command)	What the option does
<code>security.passwd.firstlogin.enable</code> {on off}	<p>Specifies whether the password must be changed when new users log in for the first time or when users try to log in after their password has been changed by an administrator.</p> <p>The default value for this option is <code>off</code>.</p> <p>Note: If you enable this option, you must ensure that all groups have the <code>login-telnet</code> and <code>cli-passwd*</code> capabilities. Users in groups that do not have these capabilities cannot log in to the storage system.</p>

Password rule option (used with the options command)	What the option does
<code>security.passwd.lockout.numtries num</code>	<p>Specifies the number of allowable login attempts before a nonroot user's account is disabled.</p> <p>The default value for this option is 4,294,967,295.</p>
<code>security.passwd.rules.enable {on off}</code>	<p>Specifies whether a check for password composition is performed when new passwords are specified.</p> <p>If this option is set to <code>on</code>, passwords are checked against the rules specified in this table, and the password is rejected if it does not pass the check.</p> <p>If this option is set to <code>off</code>, the check is not performed.</p> <p>The default value for this option is <code>on</code>.</p> <p>This option does not apply to the users <code>root</code> or <code>Administrator</code> (the NT Administrator account) if <code>security.passwd.rules.everyone</code> is set to <code>off</code>.</p>
<code>security.passwd.rules.everyone {on off}</code>	<p>Specifies whether a check for password composition is performed for all users, including the users <code>root</code> and <code>Administrator</code>.</p> <p>If this option is set to <code>off</code>, the checks do not apply to <code>root</code> or <code>Administrator</code>. The checks still apply to all other users unless the <code>security.passwd.rules.enable</code> option is also set to <code>off</code>.</p> <p>For storage systems shipped with Data ONTAP 8.0 or later, the default value for this option is <code>on</code>.</p> <p>For storage systems upgraded from a release earlier than Data ONTAP 8.0, the setting for this option stays the same as before the upgrade.</p>
<code>security.passwd.rules.history num</code>	<p>Specifies the number of previous passwords that are checked against a new password to disallow repeats.</p> <p>For storage systems shipped with Data ONTAP 8.0 or later, the default value for this option is 6. In this case, the password cannot be the same as any of the last six passwords.</p> <p>For storage systems upgraded from a release earlier than Data ONTAP 8.0, the setting for this option stays the same as before the upgrade.</p> <p>If the <code>security.passwd.rules.enable</code> option is set to <code>off</code>, this option is ignored.</p>

Password rule option (used with the options command)	What the option does
<pre>security.passwd.rules.maximum max_num</pre>	<p>Specifies the maximum number of characters a password can have.</p> <p>The default value for this option is 256.</p> <p>Note:</p> <p>This option can be set to a value greater than 16, but a maximum of 16 characters are used to match the password.</p> <p>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.</p> <p>If the <code>security.passwd.rules.enable</code> option is set to <code>off</code>, this option is ignored.</p>
<pre>security.passwd.rules.minimum min_num</pre>	<p>Specifies the minimum number of characters a password must have.</p> <p>The default value for this option is 8.</p> <p>If the <code>security.passwd.rules.enable</code> option is set to <code>off</code>, this option is ignored.</p>
<pre>security.passwd.rules.minimum.alphabetic min_num</pre>	<p>Specifies the minimum number of alphabetic characters a password must have.</p> <p>The default value for this option is 2.</p> <p>If the <code>security.passwd.rules.enable</code> option is set to <code>off</code>, this option is ignored.</p>
<pre>security.passwd.rules.minimum.digit min_num</pre>	<p>Specifies the minimum number of digit characters a password must have. These are numbers from 0 to 9.</p> <p>The default value for this option is 1.</p> <p>If the <code>security.passwd.rules.enable</code> option is set to <code>off</code>, this option is ignored.</p>
<pre>security.passwd.rules.minimum.symbol min_num</pre>	<p>Specifies the minimum number of symbol characters (including white space and punctuation characters) a password must have.</p> <p>The default value for this option is 0.</p> <p>If the <code>security.passwd.rules.enable</code> option is set to <code>off</code>, this option is ignored.</p>

The diagnostic account and the systemshell

A diagnostic account, named “diag,” is provided with your storage system. You can enable the diagnostic account to perform troubleshooting tasks in the systemshell. The diagnostic account and the systemshell are intended only for low-level diagnostic purposes and should be used only with guidance from technical support.

The diagnostic account is the only account that can be used to access the systemshell, through the advanced command `systemshell`. The diagnostic account is disabled by default. You must enable the account and set up its password before using it. Neither the diagnostic account nor the systemshell is intended for general administrative purposes.

Next topics

[Enabling and disabling the diagnostic account](#) on page 146

[Setting the password for the diagnostic account](#) on page 147

[Accessing the systemshell](#) on page 148

Enabling and disabling the diagnostic account

With guidance from technical support, you can enable the diagnostic account to gain access to the systemshell to perform low-level diagnostic and troubleshooting tasks. You can also disable the diagnostic account at any time to disallow access to the systemshell.

Steps

1. Set your privilege level to advanced by entering the following command at the storage system prompt:

```
priv set advanced
```

2. Do one of the following:

If you want to...	Enter the following command at the storage system prompt...
Display the diagnostic account information and status	<pre>useradmin diaguser show</pre> <p>By default, the diagnostic account is disabled.</p> <p>Note: The diagnostic account user name, “diag,” is not displayed as a part of the <code>useradmin user list</code> command. To display the account information, you must use <code>useradmin diaguser show</code>.</p>
Enable the diagnostic account.	<pre>useradmin diaguser unlock</pre>
Disable the diagnostic account	<pre>useradmin diaguser lock</pre>

Example of the `useradmin diaguser` command output

The following example shows how you can use the `useradmin diaguser` commands to display and enable the diagnostic account.

```
nodename*>useradmin diaguser show
Name: diag
Info: Account for access to systemshell
Locked: yes

nodename*>useradmin diaguser unlock

nodename*>useradmin diaguser show
Name: diag
Info Account for access to systemshell
Locked: no
```

Setting the password for the diagnostic account

After enabling the diagnostic account, you must set the password for the account before you can use it to access the systemshell.

Steps

1. Set your privilege level to advanced by entering the following command at the storage system prompt:

```
priv set advanced
```

2. Enter the following command at the storage system prompt to set the password for the diagnostic account:

```
useradmin diaguser password
```

The following password rules apply to the diagnostic account:

- The password cannot contain the user name.
- The password must be at least eight characters long.
- The password must contain at least one letter and one number.
- The password cannot be the same as the last six passwords.

Example of the `useradmin diaguser password` command output

The following example shows how you can use the `useradmin diaguser password` command to set the password for the diagnostic account.

```
nodename*>useradmin diaguser password
Please enter a new password:
```

```
Please enter it again:
```

Accessing the systemshell

The systemshell is intended only for low-level diagnostic purposes.

Before you begin

Only the diagnostic account user, named “diag,” can access the systemshell. Before accessing the systemshell, ensure that the diagnostic account has been enabled (using `useradmin diaguser unlock`) and the password has been set (using `useradmin diaguser password`).

About this task

The systemshell is not intended for general administrative purposes and should only be used with guidance from technical support. Misuse of the systemshell can result in system failure and data loss or corruption.

Steps

1. If necessary, change the privilege level to advanced by entering the following command at the storage system prompt:

```
priv set advanced
```

2. Enter the following command to enter the systemshell:

```
systemshell
```

This command takes no arguments and invokes the diagnostic account login.

Note: If the diagnostic account is disabled or the password is not set, attempts to log in to the systemshell will fail.

3. To exit the systemshell and return to the storage system prompt, enter the following command:

```
exit
```

Example of the `systemshell` command output

The following example shows the screen output of the `systemshell` command when the diagnostic account has been enabled and the password has been set.

```
nodename*>systemshell
Data ONTAP/i386 (nodename) (tty0)
login: diag
Password:
Last login: Thu Mar 26 19:35:55 from localhost
WARNING: The systemshell provides access to low-level
```

```
diagnostic tools that can cause irreparable damage to
the system if not used properly. Use this environment
only when directed to do so by support personnel.
```

```
%whoami
diag

%exit
logout

nodename*>
```

The following example shows the screen output of the `systemshell` command when the diagnostic account is disabled.

```
nodename*>useradmin diaguser lock

nodename*>useradmin diaguser show
Name: diag
Info: Account for access to systemshell
Locked: yes

nodename*>systemshell

Data ONTAP/i386 (nodename) (tty0)

login: diag
Password:
Login incorrect
login: diag
Password:
Login incorrect
(CTRL-C)

nodename*>
```

Related tasks

[Enabling and disabling the diagnostic account](#) on page 146

[Setting the password for the diagnostic account](#) on page 147

General system maintenance

General maintenance tasks you might need to perform to manage your storage system include managing aggregate Snapshot copy; managing licenses; setting the system date and time; synchronizing the system time; managing core files; configuring message logging, audit logging, and storage system startup; backing up and cloning storage system configuration; and managing UPS.

Next topics

- [Special system files](#) on page 151
- [Aggregate Snapshot copy management](#) on page 151
- [Ways to manage licenses](#) on page 154
- [Setting the system date and time](#) on page 156
- [Synchronizing the system time](#) on page 157
- [Displaying and setting the system time zone](#) on page 158
- [Core files](#) on page 159
- [Message logging](#) on page 160
- [Audit logging](#) on page 163
- [Startup configuration for the storage system](#) on page 165
- [Storage system configuration backup and cloning](#) on page 168
- [About writing and reading files on the storage system](#) on page 170
- [UPS management](#) on page 172

Special system files

For storage systems upgraded from a release earlier than Data ONTAP 8.0, some system files exist in every volume of the system. You must not remove or modify these files unless technical support directs you to do so. These files enable you to restore LUNs in Snapshot copies if you revert to a release earlier than Data ONTAP 8.0.

The following system files are in the root level of every volume, including the root volume:

- `.vtoc_internal`
- `.bplusvtoc_internal`

Aggregate Snapshot copy management

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. 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 the mirror, an aggregate Snapshot copy is created automatically before 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.

For more information about Snapshot copies, see the *Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide*.

Next topics

[How to create aggregate Snapshot copies](#) on page 152

[Aggregate Snapshot reserve](#) on page 152

[Automatic aggregate Snapshot copy deletion](#) on page 153

[Disabling automatic aggregate Snapshot copy creation](#) on page 154

How to 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 (`snap create`) as you would for a volume Snapshot copy, except that you add the `-A` flag. For more information about creating and managing Snapshot copies, see the *Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide* and the `na_snap(1)` man page.

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. Usually, the default aggregate Snapshot reserve of 5 percent is sufficient. However, you might increase the aggregate Snapshot reserve under some circumstances.

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

Note: Aggregate Snapshot copies cannot consume any space outside of their Snapshot reserve, if automatic aggregate Snapshot copy deletion is enabled. If automatic aggregate Snapshot copy deletion is disabled, then aggregate Snapshot copies can consume space outside of their Snapshot reserve.

You should consider increasing the aggregate Snapshot reserve if:

- You find that aggregate Snapshot copies are being created and deleted often enough to affect system performance.
- You need to complete a mirror resync when an aggregate is being written to very frequently. In this case, the default aggregate Snapshot reserve may not be large enough to hold all the resync Snapshot copies until the resync completes.

For information about how your system is using space reserved for aggregates, including aggregate Snapshot copies, use the `aggr show_space` command. See the `na_aggr(1)` man page for detailed information.

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 5 percent. If you need to reclaim the space being used for the aggregate Snapshot reserve, disable automatic aggregate Snapshot copy creation.

Related tasks

[Disabling automatic aggregate Snapshot copy creation](#) on page 154

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 [waf1.snap.autoDelete:info]: Deleting snapshot 'nightly.0' in aggregate 'aggr1' to recover storage

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 one plex of a RAID SyncMirror aggregate has to be offline for some time, you would want to make sure that the SyncMirror-based Snapshot copy is not automatically deleted.

To disable automatic aggregate Snapshot copy deletion, you use the `aggr options` 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 reenable automatic aggregate Snapshot copy deletion as quickly as possible.

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. Disabling automatic aggregate Snapshot copy creation reclaims the free space used for the aggregate Snapshot reserve. However, you are advised to leave automatic aggregate Snapshot copy creation enabled, in case you need any low-level file system repair.

About this task

If you have a MetroCluster configuration or if you are using RAID SyncMirror, ensure that no creation of aggregate Snapshot copies is scheduled. If Snapshot creation has been scheduled, an error message is displayed, advising you to turn off scheduled creation of aggregate Snapshot copies to reduce the chances of running out of space for aggregate Snapshot copies.

Steps

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
```

Ways to manage licenses

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. You can add or disable a license. You can also display the licensing information for your storage system.

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

You can perform the following tasks to manage licenses:

- Add licenses

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

Next topics

[Adding a license](#) on page 155

[Displaying current license codes](#) on page 155

[Disabling a license](#) on page 155

Adding a license

If a service requires license, you must add the license code to the storage system before you can use the service.

Step

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

You can display licensing information for all services that are enabled for your storage system.

Step

1. Enter the following command without parameters:

```
license
```

Data ONTAP displays a list of the licenses that are enabled and their codes.

Disabling a license

You can disable a licensed service, making it unavailable for the storage system.

About this task

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

Step

1. Enter the following command:

```
license delete service
```

service is one of the list of possible services.

Setting the system date and time

Keeping the system date and time correct is important to ensure that the storage system can service requests correctly.

About this task

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 Snapshot copy with the earlier date was created before the one with the later date, 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.

Steps

1. Access the storage system command line through the console or through a Telnet session.
2. Enter the following command, substituting the current date and time for the number string:

```
date [-u] [[CC]yy]mmddhhmm[.ss]
```

`-u` sets the date and time to Greenwich Mean Time instead of the local time.

`CC` is the first two digits of the current year.

`yy` is the second two digits of the current year.

`mm` is the current month. If the month is omitted, the default is the current month.

`dd` is the current day. If the day is omitted, the default is the current day.

`hh` is the current hour, using a 24-hour clock.

`mm` is the current minute.

`ss` is the current second. If the seconds are omitted, the default is 0.

Example

The following command sets the date and time to 22 May 2002 at 9:25 a.m.

```
date 200205220925
```

Note: If the first two digits of the year are omitted, they default to 20; if all four digits are omitted, they default to the current year. Time changes for daylight saving and standard time, and for leap seconds and years, are handled automatically.

Synchronizing the system time

The `timed` daemon enables you to keep the system time for your storage system automatically synchronized with a time server. Using this feature is advised, because problems can occur when the storage system clock is inaccurate.

About this task

To keep your storage system time synchronized automatically, 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.

You use the Network Time Protocol (NTP) protocol for time synchronization. You can get a list of public NTP time servers from the NTP.Servers Web at <http://ntp.isc.org/bin/view/Servers/WebHome>.

Steps

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.
2. Set the appropriate `timed` options by using the `options` command at the storage system prompt.

At a minimum, you must ensure that the `timed.proto` option is set to `ntp`, and set the `timed.servers` option to at least one valid time server.

You must also ensure that the `timed.enable` option is set to `on`.

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

Related tasks

[Setting the system date and time](#) on page 156

The timed options

The `timed` options support features such as enabling time synchronization and specifying the servers to use for time synchronization.

The following table describes the `timed` options.

Option	Function	Values	Default
<code>timed.enable</code>	Enables time synchronization.	<ul style="list-style-type: none"> • <code>on</code> • <code>off</code> 	<code>on</code>
<code>timed.log</code>	Specifies whether time changes should be logged to the console.	<ul style="list-style-type: none"> • <code>on</code> • <code>off</code> 	<code>off</code>

Option	Function	Values	Default
<code>timed.proto</code>	Specifies the protocol used to synchronize the time.	<code>ntp</code>	<code>ntp</code>
<code>timed.servers</code>	Specifies up to five time servers used by the <code>timed</code> features.	For example, <code>times1, times2.company.com, 10.15.46.92</code>	null string

For more detailed information on the `timed` options, see the `na_options(1)` man page.

Example of clock synchronization

The following example configures `timed` to use the NTP protocol.

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

Displaying and setting the system time zone

Data ONTAP enables you to display the system time zone. It also enables you to set the system time zone and save the setting for use on subsequent boots.

Steps

1. Access the storage system command line through the console or through a Telnet session.
2. Enter the following command:

```
timezone [name]
```

The *name* argument specifies the time zone to use. Each time zone is described by a file in the storage system's `/etc/zoneinfo` directory. The *name* argument is the file name under `/etc/zoneinfo` that describes the time zone to use. If no argument is specified, the current time zone name is displayed.

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

Example

The following commands set the time zone to the time zone file `/etc/zoneinfo/America/Los_Angeles` and display the set time zone.

```
toaster> timezone America/Los_Angeles
toaster> timezone
Current time zone is America/Los_Angeles
```

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.

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.

Next topics

[Core dump writing](#) on page 159

[Automatic technical support notification upon system reboots](#) on page 160

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 is in a high-availability configuration and the `cf.takeover.on_panic` option 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.

The `coredump.dump.attempts` option controls how many attempts are made to create a core dump file. The default value is 2.

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

Automatic technical support notification upon system reboots

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.

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

Message logging

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.

You can access the `/etc/messages` files using your NFS or CIFS client, or using HTTP(S).

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.

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).

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.

Next topics

[The `/etc/syslog.conf` file](#) on page 161

[Sample `/etc/syslog.conf` file](#) on page 162

[Configuring message logging](#) on page 162

Related concepts

[How to access the default directories on the storage system](#) on page 97

[The `/etc/messages` file](#) on page 96

[How to access the default directories on the storage system](#) on page 97

Related tasks

[Accessing log files using HTTP or HTTPS](#) on page 100

The /etc/syslog.conf file

The `/etc/syslog.conf` file configures the level of information that the storage system records. It specifies the subsystem from which the message originated, the severity of the message, and where the message is sent.

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 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 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

Level	Description
debug	Used for diagnostic purposes
*	All levels of errors

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.	<code>/etc/messages</code>
Send messages to a host name preceded by an @ sign.	<code>@adminhost</code>
Send messages to the console.	<code>/dev/console</code> or <code>*</code>

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

Sample `/etc/syslog.conf` file

The sample 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

The `/etc/syslog.conf` file can be edited to modify your system's message logging.

Steps

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

3. Save and close the `/etc/syslog.conf` file.

The changes you made to the `syslog.conf` file are read automatically and are reflected in the message logging.

Related concepts

[The `/etc/syslog.conf` file](#) on page 161

Audit logging

An audit log is a record of commands executed at the console, through a Telnet shell or an SSH shell, or by using the `rsh` command. All the commands executed in a source file script are also recorded in the audit log. Administrative HTTP operations, such as those resulting from the use of FilerView, are logged. All login attempts to access the storage system, with success or failure, are also audit-logged.

In addition, changes made to configuration and registry files are audited. Read-only APIs by default are not audited but you can enable auditing with the `auditlog.readonly_api.enable` option.

By default, Data ONTAP is configured to save an audit log. The audit log data is stored in the `/etc/log` directory in a file called `auditlog`.

For configuration changes, the audit log shows the following information:

- What configuration files were accessed
- When the configuration files were accessed
- What has been changed in the configuration files

For commands executed through the console, a Telnet shell, an SSH shell, or by using the `rsh` command, the audit log shows the following information:

- What commands were executed
- Who executed the commands
- When the commands were executed

The maximum size of the audit-log file is specified by the `auditlog.max_file_size` option. The maximum size of an audit entry in the audit-log file is 200 characters. An audit entry is truncated to 200 characters if it exceeds the size limit.

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 audit-log file reaches the maximum size specified by `auditlog.max_file_size`.

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

You can access the audit-log files using your NFS or CIFS client, or using HTTP.

Note: You can also configure auditing specific to your file access protocol. For more information, see the *Data ONTAP 7-Mode File Access and Protocols Management Guide*.

For information about forwarding audit logs to a remote syslog log host, see the `na_auditlog(5)` man page.

Next topics

[Configuring audit logging](#) on page 164

[Enabling or disabling read-only API auditing](#) on page 164

Related concepts

[How to access the default directories on the storage system](#) on page 97

Configuring audit logging

You can change the maximum size of the audit log file.

Steps

1. If audit logging is turned off, enter the following command to turn audit logging on:

```
options auditlog.enable on
```

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).

Enabling or disabling read-only API auditing

Data ONTAP enables you to control auditing of APIs based on their roles. If an API is used only for retrieving information and not for modifying the state of the system, the read-only API is not audited by default.

About this task

You use the `auditlog.readonly_api.enable` option to enable or disable read-only API auditing. The default value of the `auditlog.readonly_api.enable` option is `off`. It is recommended that you leave this option disabled, because auditing read-only APIs may inundate the audit log.

Step

1. Enter the following command to enable or disable read-only API auditing:

```
options auditlog.readonly_api.enable {on|off}
```

The default is `off`.

Startup configuration for the storage system

You can customize your system startup by editing the storage system's boot configuration file, the `/etc/rc` file in the root directory.

Next topics

[About the `/etc/rc` file](#) on page 165

[Editing the `/etc/rc` file](#) on page 166

[Recovering from `/etc/rc` errors](#) on page 167

About the `/etc/rc` file

Startup commands for your storage system are stored in the `/etc/rc` file. The `/etc/rc` file contains commands that the storage system executes at boot time to configure the system.

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

Some commands cannot be stored in the `/etc/rc` file. This includes commands that are executed by subsystems that are not yet available when the `/etc/rc` file is executed. For example, you cannot include `iscsi` commands in the `/etc/rc` file. Doing so prevents your storage system from booting successfully.

Running the `setup` command rewrites the `/etc/rc` file. You should back up the `/etc/rc` file if you must rerun the `setup` command after the system's initial setup.

Sample `/etc/rc` file

The sample `/etc/rc` file shows default startup commands.

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
```

The following table explains the sample `/etc/rc` file:

Description	Explanation
<code>hostname toaster</code>	Sets the storage system host name to “toaster.”
<code>ifconfig e0 `hostname`-0</code> <code>ifconfig e1 `hostname`-1</code> <code>ifconfig f0 `hostname`-f0</code> <code>ifconfig a5 `hostname`-a5</code>	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 <code>/etc/hosts</code> file on the storage system root volume. If you prefer to have the actual IP addresses in the <code>/etc/rc</code> file, you can enter IP addresses directly in <code>/etc/rc</code> on the root volume.
<code>route add default MyRouterBox</code>	Specifies the default router. You can set static routes for the storage system by adding route commands to the <code>/etc/rc</code> file. The network address for <code>MyRouterBox</code> must be in <code>/etc/hosts</code> on the root volume.
<code>routed on</code>	Starts the routing daemon.
<code>savecore</code>	Saves the core file from a system panic, if any, in the <code>/etc/crash</code> directory on the root volume. Core files are created only during the first boot after a system panic.

For more information about the `ifconfig` command and routing, see the *Data ONTAP 7-Mode Network Management Guide*.

Related concepts

[Core files](#) on page 159

Editing the `/etc/rc` file

You edit the storage system's boot configuration file, the `/etc/rc` file, to modify the commands that the system runs at boot time.

About this task

The storage system's boot configuration file is named `rc` and is in the `/etc` directory of its default volume (the default is `/vol/vol0/etc/rc`).

Steps

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

The storage system can become inaccessible to the administration host due to errors. You can recover from the `/etc/rc` errors to make the system accessible again.

About this task

The following are some `/etc/rc` errors that might cause the system to become in accessible:

- 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.

Steps

1. Enter one of the following commands on the console to configure the interface with the correct address.

If you are in...	Then...
An NFS environment	Enter the <code>exportfs</code> 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

The configuration backup operation of the storage system stores the system's configuration information in a 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.

Next topics

[Backing up a storage system configuration](#) on page 168

[Cloning a storage system configuration](#) on page 169

[Restoring a storage system configuration](#) on page 169

[Comparing storage system configurations and backup configuration files](#) on page 170

Backing up a storage system configuration

When you 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.

Step

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.

Examples of `config dump` command

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
```

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

You can clone the configuration of one storage system to another system.

Step

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 of config clone command

The following is an example of the `config clone` command cloning the `tpubs-dot` configuration to the storage system `toaster`.

```
config clone tpubs-dot root:hello
```

Restoring a storage system configuration

You can restore storage system configuration information from a backup configuration file.

About this task

Illegal entries in the configuration file might cause attempts to fail and error messages to occur when using `config restore -v` to restore volume-specific configurations. If this happens, edit the configuration file in the default `/etc/configs` directory to remove the illegal entries.

For instance, an error message indicating an invalid operation on FlexVol volume `vol_name` could result from a configuration file containing the text `options.vols.vol_name.raidsize`, where `vol_name` is not a traditional volume and thus an illegal entry that should be removed from the configuration file.

Steps

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.

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

Example of `config restore` command

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
```

Comparing storage system configurations and backup configuration files

You can compare a storage system's current configuration with a backup configuration file to see the difference. You can also compare differences between two backup configuration files.

Step

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.

Examples of `config diff` command

The following example compares the storage system's current configuration with the configuration information in the backup file.

```
config diff 11_15_2004
```

The following example compares the configuration information in two backup files.

```
config diff -o diff.txt 11_05_2004 11_15_2004
```

About writing and reading files on the storage system

Data ONTAP provides commands that enable you to write to or read from a specified file on the storage system. However, when using such commands, you must exercise caution about potential security and data corruption issues.

Next topics

[Writing a WAFL file](#) on page 171

[Reading a WAFL file](#) on page 172

Writing a WAFL file

Data ONTAP enables you to read data from standard input and write it into the specified file.

About this task

A user who has the capability to execute the `wrfile` command can write over or append data to any file on the storage system. Exercise caution about security and data corruption issues when using the `wrfile` command.

Step

1. Enter the following command:

```
wrfile [-a] filename [...]
```

filename is the name of the file you want to write or append to. It must be a fully qualified path name. If *filename* does not already exist, the `wrfile` command will create it.

The `-a` option appends the rest of the command line after *filename* to the file. If the `-a` option is not used, the `wrfile` command closes the file when it reads an EOF from the input stream or, if run on the console, when interrupted by the interrupt character.

Note: There are restrictions for using the `-a` option with special characters, # (hash), ` (backtick), and " (double quotation marks). In general, if you use the `-a` option, surround the line to be written with quotation marks.

The interrupt character is Ctrl-C. If `wrfile` is run from the console, interrupting `wrfile` causes all characters typed on the same line as the interrupt character to be lost. The storage system will also issue an "interrupted system call" error message.

Example of `wrfile` command

The following example uses `wrfile` to create a file `/etc/test` that contains two lines, "line#1" and "line#2".

```
toaster> wrfile /etc/test
line#1
```

Press Enter, followed by the interrupt character (Ctrl-C).

```
read: error reading standard input: Interrupted system call
toaster> wrfile -a /etc/test "line#2"
toaster>
```

See the `na_wrfile(1)` man page for additional examples.

Related tasks

[Reading a WAFL file](#) on page 172

Reading a WAFL file

Data ONTAP enables you to read a file from the storage system and write its contents to standard output.

About this task

A user who has the capability to execute the `rdfile` command can read any file on the storage system. Exercise caution about security issues with the `rdfile` command.

Step

1. Enter the following command:

```
rdfile filename
```

filename is the name of the file whose content you want to read. It must be a fully qualified path name.

Note: Files that contain non-ASCII characters may have indeterminate output.

Example of `rdfile` command

The following example uses the `rdfile` command to read the content of the `/etc/test` file, which contains two lines, "line#1" and "#line#2".

```
toaster> rdfile /etc/test
line#1
line#2
toaster>
```

Related tasks

[Writing a WAFL file](#) on page 171

UPS management

Data ONTAP enables you to register and monitor the status of Uninterruptible Power Supply (UPS) 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.

For more information about the `ups` command, see the `na_ups(1)` man page.

Next topics

[The UPS shutdown options](#) on page 173

[The UPS shutdown process](#) on page 173

Factors that might influence UPS shutdown event timing on page 174

The UPS shutdown options

Data ONTAP provides two configurable values, `warningtime` and `criticaltime`, 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).

Note: Using the `ups set-limits` command, you can set the UPS battery critical time and warning time for all UPS devices or for a specific UPS device by specifying its IP address. You can display the UPS battery critical time and warning time by using the `ups print-limits` command. For more information, see the `na_ups(1)` man page.

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.

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

Attention: You are strongly advised to contact technical support before changing the shutdown event timing values.

The UPS 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 following is the shutdown process:

1. When a power loss occurs, an SNMP trap, AutoSupport message, and log messages are generated alerting you that the power loss has occurred.

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

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.

Factors that might influence UPS shutdown event timing

The factors that can affect shutdown event timing include the UPS battery availability, the storage system workload, and your company policies and procedures

- 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.

The AutoSupport feature

AutoSupport 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.

Next topics

[Overview of the AutoSupport feature](#) on page 175

[Configuring AutoSupport](#) on page 177

[AutoSupport options](#) on page 177

[Testing AutoSupport](#) on page 182

[AutoSupport troubleshooting tasks](#) on page 183

[AutoSupport messages](#) on page 185

Overview of the AutoSupport feature

The AutoSupport feature 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` feature is enabled by default on the storage system.
AutoSupport is enabled by default when you configure your storage system for the first time. AutoSupport begins sending messages to technical support 24 hours after AutoSupport is enabled. You can cut short the 24-hour period by upgrading or reverting the system, modifying the AutoSupport configuration, or changing the time of the system to be outside of the 24-hour period.
Note: You can disable AutoSupport at any time by turning off the `autosupport.enable` option, but you should leave it enabled. Enabling AutoSupport can significantly help speed problem determination and resolution should a problem occur on your storage system.
Although AutoSupport messages to technical support are enabled by default, you need to set the correct options and have a valid mail host to have messages sent to your internal support organization.
- 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, between 12:00 a.m. and 1 a.m. Sunday
Three AutoSupport messages are generated during this time: The first, the weekly AutoSupport message, provides the same system information as regular AutoSupport messages. The second, 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 third, the NetApp Health Trigger (NHT) message, provides information about any failed disk drives. If no drives failed during the past week, no weekly drive NHT message is sent. By default, the drive NHT message is sent only to technical support.
- The system can send AutoSupport messages by SMTP, HTTP, or HTTPS (Hypertext Transfer Protocol over Secure Sockets Layer). HTTPS is the default, and you should use it whenever possible.
- If an AutoSupport message cannot be sent successfully, an SNMP trap is generated.

For more information about AutoSupport, see the NOW site.

Related information

<http://now.netapp.com/>

AutoSupport transport protocols

AutoSupport supports HTTPS, HTTP, and SMTP as the transport protocols for delivering AutoSupport messages to technical support. If you enable AutoSupport messages to your internal support organization, those messages are sent by SMTP.

AutoSupport supports the following types of transport protocols:

- HTTPS (This is the default transport protocol used by AutoSupport; you should use it whenever possible.)
- HTTP
- SMTP

Because SMTP can introduce limitations on message length and line length, you should use HTTP or, preferably, HTTPS for your AutoSupport transport protocol if possible.

HTTP uses port 80; HTTPS uses port 443. If the network connection does not allow HTTPS or HTTP, you need to configure AutoSupport for SMTP. SMTP uses port 25.

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

The AutoSupport feature 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 mail hosts in the AutoSupport e-mail delivery pathway are configured to send and receive the 8-bit Multipurpose Internet Mail Extensions (MIME) encoding.

Configuring AutoSupport

To configure AutoSupport, you specify AutoSupport options.

About this task

Modifications to AutoSupport options are persistent across reboots except for the `autosupport.doit`, `autosupport.performance_data.doit`, `autosupport.support.to`, and `autosupport.support.url` options.

Step

1. Enter the following command:

```
options autosupport.option [arguments]
```

option is the AutoSupport option you want to configure.

arguments is the required or optional argument for *option*.

Related concepts

[AutoSupport options](#) on page 177

AutoSupport options

You use the AutoSupport options to configure the AutoSupport feature.

The main AutoSupport options are shown in the following table. For more information, see the `na_options(1)` and the `na_autosupport(8)` man pages.

AutoSupport option	Description
<code>autosupport.cifs.verbose</code> [on off]	Enables and disables inclusion of CIFS session and share information in AutoSupport messages. The default is <code>off</code> .

AutoSupport option	Description
<code>autosupport.content</code> [complete minimal]	<p>Indicates the type of content that AutoSupport messages should contain. The default is <code>complete</code>.</p> <p>Note: You should keep the setting at <code>complete</code>. Changing the setting to <code>minimal</code> limits the ability of technical support to respond quickly to problems.</p> <p>If you change from <code>complete</code> to <code>minimal</code>, any AutoSupport message with <code>complete</code> content not yet sent is cleared from the outgoing message spool and a message to that effect appears on the console.</p>
<code>autosupport.doit</code> [<i>message</i>]	<p>Tells the <code>autosupport</code> feature to send an AutoSupport notification immediately.</p> <p>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.</p> <p>You can verify that AutoSupport is working by using the “Call Home Check” function, which sends an <code>autosupport.doit</code> message with a subject line containing any variation of the word <code>TEST</code> or <code>TESTING</code>. 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.</p>
<code>autosupport.enable</code> [on off]	<p>Enables and disables AutoSupport notification. The default is <code>on</code>.</p>
<code>autosupport.from</code> <i>sender</i>	<p>Defines the user to be designated as the sender of the notification, for example, <code>postmaster@mycompany.com</code>.</p> <p>The default is “Postmaster@xxx” where xxx is the name of the system.</p>
<code>autosupport.local.nht_data.enable</code>	<p>Sends a copy of NetApp Health Trigger (NHT) disk data to the customer “to” list. The default is <code>off</code>.</p> <p>For this option to work, you also need to specify a mail host with the <code>autosupport.mailhost</code> option and an address with the <code>autosupport.to</code> option.</p>

AutoSupport option	Description
<code>autosupport.local.performance_data.enable</code>	<p>Determines whether the weekly performance AutoSupport message is sent to all of the recipients designated by the <code>autosupport.to</code> option or only to technical support. The default is <code>off</code>.</p> <p>For this option to work, you also need to specify a mail host with the <code>autosupport.mailhost</code> option and an address with the <code>autosupport.to</code> option.</p>
<code>autosupport.mailhost host1[, ..., host5]</code>	<p>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 "mailhost."</p> <p>The specified mail hosts will be used to send AutoSupport messages.</p>
<code>autosupport.minimal.subject.id [hostname systemid]</code>	<p>Defines how the system is identified in the AutoSupport message title if <code>autosupport.content</code> is <code>minimal</code>. The default is <code>hostname</code>.</p>
<code>autosupport.nht_data.enable</code>	<p>Sends a copy of NetApp Health Trigger (NHT) disk drive data to technical support. The default is <code>on</code>.</p>
<code>autosupport.noteto address1[, ..., address5]</code>	<p>Defines the list of recipients for the AutoSupport short note e-mail. The short note e-mail consists only of the subject line of the AutoSupport message, which is easily viewed on a cell phone or other text device.</p> <p>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.</p> <p>You can have AutoSupport messages sent to your internal support organization by setting this option (or the <code>autosupport.to</code> option) and having a valid mail host.</p>
<code>autosupport.partner.to address1[, ..., address5]</code>	<p>Defines the list of recipients who will receive all AutoSupport e-mail notifications regardless of the severity level.</p> <p>Up to five e-mail addresses are allowed. Enter the addresses as a comma-separated list with no spaces in between. By default, no list is defined.</p> <p>This option is not affected by the setting of the <code>autosupport.support.enable</code> option.</p>
<code>autosupport.performance_data.doit any_string</code>	<p>Triggers a performance snapshot AutoSupport message when any string is added.</p>

AutoSupport option	Description
<code>autosupport.performance_data.enable</code>	Enables the weekly performance AutoSupport messages to technical support. This option should always be set to <code>on</code> . If you do not want the weekly performance AutoSupport message to be sent to all recipients on the list defined in the <code>autosupport.to</code> option, disable the <code>autosupport.local.performance_data.enable</code> option. The default is <code>on</code> .
<code>autosupport.periodic.tx_window <i>time</i></code>	Specifies the randomized delay window for periodic AutoSupport messages. Values can range from 0 seconds to 240 (4 hours). The default is 60 (1 hour). Setting the value to 0 disables the randomized delay, which is intended to prevent bottlenecks.
<code>autosupport.retry.count <i>#retries</i></code>	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 4,294,967,294. The default is 15.
<code>autosupport.retry.interval <i>interval</i></code>	Defines the time to wait before trying to resend a failed AutoSupport notification. The values can end with <code>s</code> , <code>m</code> , or <code>h</code> 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).
<code>autosupport.support.enable [on off]</code>	Enables and disables sending AutoSupport notification to technical support. The default is <code>on</code> .

AutoSupport option	Description
<pre>autosupport.support.proxy [user:pass@]proxyhost.com[:port][/]</pre>	<p>Allows you to set an HTTP proxy if necessary. This is useful only if <code>autosupport.support.transport</code> is set to <code>http</code> or <code>https</code>. The default value for this option is an empty string.</p> <p>You use this option to specify user name and password for proxy authentication. The URL is entered without an <code>http://</code> or <code>https://</code> prefix. The following are some examples:</p> <ul style="list-style-type: none"> • <code>options autosupport.support.proxy myusername:mypassword@myhost.com</code> • <code>options autosupport.support.proxy myusername:mypassword@myhost.com:9090</code> • <code>options autosupport.support.proxy myhost.com</code> • <code>options autosupport.support.proxy myhost.com:9090</code> <p>Note: The value you use for this option is site-specific; see your IT department for the correct value for your site.</p> <p>Note: Proxy configuration defaults to port 3128 when no port is specified.</p>
<pre>autosupport.support.to</pre>	<p>Indicates where AutoSupport notifications are sent if <code>autosupport.support.transport</code> is <code>smtp</code>. This option is read-only and is shown for informational purposes only.</p>
<pre>autosupport.support.transport [http https smtp]</pre>	<p>Defines the type of delivery for AutoSupport notifications. The default is <code>https</code>.</p>
<pre>autosupport.support.url</pre>	<p>Indicates where AutoSupport notifications are sent if <code>autosupport.support.transport</code> is <code>http</code> or <code>https</code>. This option is read-only and is shown for informational purposes only.</p>
<pre>autosupport.throttle [on off]</pre>	<p>Drops additional messages when too many AutoSupport messages of the same type are sent in too short a time. The default is <code>on</code>.</p>

AutoSupport option	Description
<code>autosupport.to address1[, ..., address5]</code>	<p>Defines the list of recipients for the AutoSupport e-mail notification. Recipients defined in this option receive only critical AutoSupport e-mail notifications; however, all AutoSupport notifications, regardless of their level of severity, continue to be sent to technical support as displayed by the read-only option <code>autosupport.support.to</code>.</p> <p>Up to five e-mail addresses are allowed, or the list can be left empty.</p> <p>Enter the addresses as a comma-separated list with no spaces in between. The default is no list.</p> <p>The addresses should include your system administrator or administrative group.</p> <p>You can have AutoSupport messages sent to your internal support organization by setting this option (or the <code>autosupport.noteto</code> option) and having a valid mail host.</p>

Related concepts

[AutoSupport troubleshooting tasks](#) on page 183

Testing AutoSupport

Testing AutoSupport helps you ensure that AutoSupport is properly configured.

Step

1. Enter the following command:

```
options autosupport.doit message
```

message is the subject line for the test AutoSupport e-mail.

If you use the keyword TEST in the message, you receive a return message indicating that the AutoSupport process is working correctly.

Related concepts

[AutoSupport options](#) on page 177

AutoSupport troubleshooting tasks

If the AutoSupport test message is not being sent, you perform the troubleshooting task to try to resolve the problem. The troubleshooting task you perform depends on the AutoSupport transport protocols you use.

Next topics

[Troubleshooting AutoSupport over HTTP or HTTPS](#) on page 183

[Troubleshooting AutoSupport over SMTP](#) on page 183

[Controlling the size of AutoSupport messages](#) on page 184

Troubleshooting AutoSupport over HTTP or HTTPS

If the AutoSupport test message is not being sent and you are using HTTP or HTTPS, check that DNS is enabled and configured correctly and that the system is routing out to the Internet successfully.

Steps

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, check that the mail host specified is a host that the storage system can talk to and that the host can serve SMTP requests.

Steps

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 `autosupport.doit` option.

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.

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
```

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
```

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.

Controlling the size of AutoSupport messages

You should control the size of AutoSupport messages. Doing so can prevent problems, especially if you use the SMTP transport protocol.

About this task

AutoSupport messages contain information about the system and the contents of various `/etc` directory files. Your e-mail server might have size limits that can cause messages to be dropped when SMTP is used as the protocol.

Steps

1. To control the size of AutoSupport messages, take one or more of the following actions:
 - Make sure the `/etc/messages` file is being rotated on a weekly basis as expected. If necessary, rotate the file manually.
 - Make sure you have your `/etc/syslog.conf` file 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.
 - Consider using HTTP or HTTPS for your AutoSupport transport protocol.

2. If these steps do not resolve the problem, you can set the `autosupport.content` option to `minimal`.

Using the `minimal` setting is not advised, because it can affect the quality of your technical support.

AutoSupport messages

AutoSupport messages help you understand the status and operations of your storage system. The AutoSupport message includes a log level that indicates the priority assignment from technical support.

The log level that indicates the priority assignment can be one of the following:

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

If you are using AutoSupport locally, you will see the log levels in the subject lines of the AutoSupport e-mail that you receive.

Next topics

[Getting AutoSupport message descriptions](#) on page 185

[Contents of AutoSupport event messages](#) on page 186

Getting AutoSupport message descriptions

The descriptions of the AutoSupport messages that you receive are available through the online Message Matrices page.

Steps

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**.

All AutoSupport message descriptions are listed alphabetically by subject line.

Related information

<http://now.netapp.com/>

Contents of AutoSupport event messages

AutoSupport messages contain various kinds of information, such as dates, version numbers, and serial numbers.

Each AutoSupport message contains the following types of information.

Note: Items in the following list 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.

- 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 commands that provide system information

Note: If you use parsers to monitor AutoSupport messages, be aware that the output of some commands might have changed since the last release of Data ONTAP and might change in future releases. For information about individual commands, see the Data ONTAP man pages.

- Checksum status
- Error-Correcting Code (ECC) memory scrubber statistics
- The following information, if HA pair configuration is licensed:**
 - System ID of the partner in an HA pair
 - Host name of the partner in an HA pair
 - HA pair node status, including the HA pair monitor and HA pair interconnect statistics
- Contents of selected `/etc` directory files
- Expiry date of all SnapLock volumes on the system*
- Registry information
- Usage information*
- Service statistics
- Boot time statistics*
- NVLOG statistics*
- WAFL check log
- Modified configurations
- X-header information
- Information about the boot device (such as the CompactFlash card)

In addition, the contents of the `/etc/messages` and `/etc/log/ems` files are sent with each AutoSupport message as .gz attachments.

You can specify the value of the `autosupport.content` option 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.

Managing storage systems remotely

You can manage your storage system remotely by using a remote management device, which can be the Service Processor (SP), the Remote LAN Module (RLM), or the Baseboard Management Controller (BMC), depending on the storage system model. The remote management device stays operational regardless of the operating state of the storage system. It provides remote platform management capabilities, including remote access, monitoring, troubleshooting, logging, and alerting features.

The SP is included in the 32xx and 62xx systems.

The RLM is included in the 30xx, 31xx, and 60xx storage system.

The BMC is included in the 20xx storage systems.

Additionally, the Remote Support Agent (RSA) is available for download as a firmware upgrade to the SP and the RLM.

Next topics

[Using the Service Processor for remote system management](#) on page 189

[Using the Remote LAN Module for remote system management](#) on page 216

[Using the Baseboard Management Controller for remote system management](#) on page 245

[The Remote Support Agent as a firmware upgrade](#) on page 273

Using the Service Processor for remote system management

The Service Processor (SP) is a remote management device that is included in the 32xx and 62xx systems. It enables you to access, monitor, and troubleshoot the storage system remotely.

The SP provides the following capabilities:

- The SP enables you to access the storage system remotely to diagnose, shut down, power-cycle, or reboot the system, regardless of the state of the storage controller.

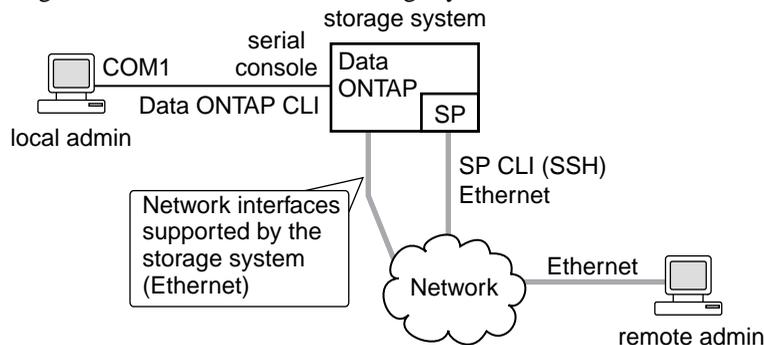
The SP is powered by a standby voltage, which is available as long as the system has input power to at least one of the system's power supplies.

The SP is connected to the system through the serial console. You can log in to the SP by using a Secure Shell client application from an administration host. You can then use the SP CLI to monitor and troubleshoot the system remotely. In addition, you can use the SP to access the system console and run Data ONTAP commands remotely.

You can access the SP from the system console or access the system console from the SP. The SP allows you to open both an SP CLI session and a separate system console session simultaneously.

- The SP monitors environmental sensors and logs system events to help you take timely and effective service actions in the event that a system problem occurs.
The SP monitors the system temperatures, voltages, currents, and fan speeds. When the SP detects that an environmental sensor has reached an abnormal condition, it logs the abnormal readings, notifies Data ONTAP of the issue, and takes proactive actions as necessary to send alerts and “down system” notifications through an AutoSupport message.
If SNMP is enabled for the SP, the SP generates SNMP traps to configured trap hosts for all “down system” events.
The SP also logs system events such as boot progress, Field Replaceable Unit (FRU) changes, Data ONTAP-generated events, and SP command history.
- Hardware-assisted takeover is available on systems that support the SP and have the SP configured.
For more information about hardware-assisted takeover, see the *Data ONTAP 7-Mode High-Availability Configuration Guide*.

The following diagram illustrates access to the storage system and the SP.



Next topics

[Ways to configure the SP](#) on page 191

[Prerequisites for configuring the SP](#) on page 191

[Configuring the SP](#) on page 192

[Accounts that can access the SP](#) on page 194

[Logging in to the SP from an administration host](#) on page 195

[Accessing the SP from the system console](#) on page 196

[SP CLI and system console sessions](#) on page 197

[How to use the SP CLI](#) on page 197

[How to use Data ONTAP to manage the SP](#) on page 203

[How the SP sensors help you monitor system components](#) on page 206

[SP commands for troubleshooting the storage system](#) on page 211

[System event log and the SP](#) on page 212

[Console log and the SP](#) on page 213

AutoSupport messages for systems with the SP on page 213

How to update the SP firmware on page 214

Troubleshooting SP connection problems on page 215

Related concepts

The eOM interface on page 45

Ways to configure the SP

Before using the SP, you must configure it for your storage system and network. You can configure the SP when you set up a new storage system. You can also configure the SP by running the `setup` or the `sp setup` command.

On a storage system that comes with the SP, you can configure the SP by using one of the following methods:

- **Initializing a new storage system**
When you power on a storage system for the first time, the `setup` command begins to run automatically. When the storage system setup process is complete, the `sp setup` command runs automatically and prompts you for SP configuration information. For more information about the system setup process, see the *Data ONTAP 7-Mode Software Setup Guide*.
- **Running the Data ONTAP `setup` command**
If you want to change both system setup and SP configuration, you use the `setup` command. The system setup process ends by initiating the `sp setup` command.
- **Running the Data ONTAP `sp setup` command directly**
If the storage system has been set up and you want to reconfigure only the SP, you can use the `sp setup` command, which omits system setup and prompts you directly for SP configuration information.

Prerequisites for configuring the SP

You need information about your network and AutoSupport settings when you configure the SP.

You need the following information:

- **Network information**
If you are using a static IP address for the SP, it must be IPv4 format and you need the following information:
 - An available static IP address for the SP
 - The netmask of your network
 - The gateway IP of your network

For information about network interfaces and management, see the *Data ONTAP 7-Mode Network Management Guide*.

- AutoSupport information

The SP sends event notifications based on the settings of the following AutoSupport options:

- `autosupport.to`
- `autosupport.mailhost`

At the minimum, consider setting the `autosupport.to` option before configuring the SP. Data ONTAP automatically sends AutoSupport configuration to the SP, allowing the SP to send alerts and notifications through an AutoSupport message. You are prompted to enter the name or the IP address of the AutoSupport mail host when you configure the SP.

Note: The SP does not rely on the storage system's `autosupport.support.transport` option to send notifications. The SP uses the Simple Mail Transport Protocol (SMTP).

Related tasks

[Configuring AutoSupport](#) on page 177

Configuring the SP

You can use the `setup` command or the `sp setup` command to configure the SP, depending on whether you want to change the system setup besides configuring the SP. You can configure the SP to use either a static or a DHCP address.

About this task

Data ONTAP supports IPv4 for static addressing. If no static or DHCP address is configured, the SP has no network connectivity.

Steps

1. At the storage system prompt, enter one of the following commands:
 - **setup**
If you want to change both system setup and SP configuration, you use the `setup` command. When the storage system setup process is complete, the `sp setup` command runs automatically and prompts you for SP configuration information.
For information about system setup, see the *Data ONTAP 7-Mode Software Setup Guide*.
 - **sp setup**
If the storage system has been set up and you want to configure only the SP, you use the `sp setup` command, which omits system setup and prompts you directly for SP configuration information.
2. When the SP setup asks you whether to configure the SP, enter **y**.
3. Enter one of the following when the SP setup asks you whether to enable DHCP on the SP:
 - To use DHCP addressing, enter **y**.
 - To use static addressing, enter **n**.

4. If you do not enable DHCP for the SP, provide the following static IP information when the SP setup prompts you to:
 - The IP address for the SP
 - The netmask for the SP
 - The IP address for the SP gateway
 - The name or IP address of the mail host to use for AutoSupport (if you use the `setup` command.)
5. At the storage system prompt, enter the following command to verify that the SP network configuration is correct:

```
sp status
```

6. At the storage system prompt, enter the following command to verify that the SP AutoSupport function is working properly:

```
sp test autosupport
```

Note: The SP uses the same mail host information that Data ONTAP uses for AutoSupport. The `sp test autosupport` command requires that you set up the `autosupport.to` option properly.

The following message is a sample of the output Data ONTAP displays:

```
Sending email messages via SMTP server at mailhost@companyname.com. If
autosupport.enable is on, then each email address in autosupport.to
should receive the test message shortly.
```

Examples of configuring the SP and displaying the configuration information

The following example shows that the SP is configured to use a static IP address:

```
toaster> sp setup
The Service Processor (SP) provides remote management capabilities
including console redirection, logging and power control.
It also extends autosupport by sending
additional system event alerts. Your autosupport settings are use
for sending these alerts via email over the SP LAN interface.
Would you like to configure the SP? y
Would you like to enable DHCP on the SP LAN interface? n
Please enter the IP address of the SP []: 192.168.123.98
Please enter the netmask of the SP []: 255.255.255.0
Please enter the IP address for the SP gateway []: 192.168.123.1
Verifying mailhost settings for SP use...
```

The following example shows that the SP is configured to use DHCP:

```
toaster> sp setup
The Service Processor (SP) provides remote management capabilities
including console redirection, logging and power control.
It also extends autosupport by sending
additional system event alerts. Your autosupport settings are use
for sending these alerts via email over the SP LAN interface.
Would you like to configure the SP? y
```

```
Would you like to enable DHCP on the SP LAN interface? y
Verifying mailhost settings for SP use...
```

The following example displays the SP status and configuration information:

```
toaster> sp status
  Service Processor      Status: Online
    Firmware Version:   1.0
    Mgmt MAC Address:    00:A0:98:01:7D:5B
    Ethernet Link:      up
    Using DHCP:         no
  IPv4 configuration:
    IP Address:         192.168.123.98
    Netmask:            255.255.255.0
    Gateway:           192.168.123.1
```

Related concepts

The AutoSupport feature on page 175

Related references

Prerequisites for configuring the SP on page 191

Accounts that can access the SP

The SP comes with an account named `naroot`. Only the SP `naroot` account and Data ONTAP user accounts with the credentials of the `admin` role or a role with the `login-sp` capability can log in to the SP. These users have access to all commands available on the SP.

For enhanced security, the SP does not allow you to log in with the Data ONTAP account name `root`. Instead, it maps the Data ONTAP root account to the SP `naroot` account. You use the SP `naroot` account and the Data ONTAP root password to log into the SP.

Note: If you disable the root account's access to the storage system, the SP `naroot` account's access to the storage system is automatically disabled.

You cannot create user accounts directly from the SP. However, users created in Data ONTAP with the credentials of the `admin` role or a role with the `login-sp` capability can log in to the SP. Changes to user account credentials on the storage system are automatically updated to the SP.

You cannot use the following generic names as account names to access the SP. Therefore, it is best not to use them as Data ONTAP account names or assign them to Data ONTAP groups that have the `admin` role or a role that includes the `login-sp` capability.

- `adm`
- `bin`
- `cli`
- `daemon`
- `ftp`

- games
- halt
- lp
- mail
- man
- netapp
- news
- nobody
- operator
- shutdown
- sshd
- sync
- sys
- uucp
- www

Related concepts

[How to manage administrator and diagnostic access](#) on page 115

[Predefined roles](#) on page 128

[Supported capability types](#) on page 129

Related tasks

[Creating a new role and assigning capabilities to roles](#) on page 131

[Modifying an existing role or its capabilities](#) on page 132

[Disabling root access to the storage system](#) on page 119

Logging in to the SP from an administration host

You can log in to the SP from an administration host to perform administrative tasks remotely, if the host has a Secure Shell client application that supports SSHv2 and you have administrative privileges on the storage system.

Before you begin

The following are the prerequisites for logging in to the SP:

- The administration host you use to access the SP must support SSHv2.
The SP does not support Telnet or RSH. The `telnet.enable` and `rsh.enable` options, which enable or disable Telnet and RSH respectively, have no effect on the SP.
- You must have access to the SP `naroot` account or a Data ONTAP user account with the credentials of the `admin` role or a role with the `login-sp` capability.

About this task

If you configured the SP to use a static IP address, and if five SSH login attempts from an administration host fail consecutively within 10 minutes, the SP rejects SSH login requests and

suspends the communication with the IP address of the administration host for 15 minutes. The communication resumes after 15 minutes, and you can try to log in to the SP again.

The SP ignores the `autologout.telnet.timeout` and `autologout.console.timeout` options. The settings for these options do not have any effect on the SP.

Steps

1. Enter the following command from the administration host to log in to the SP:

```
ssh username@SP_IP_address
```

2. If you are prompted, enter the password for `username`.

The SP prompt appears, indicating that you have access to the SP CLI.

Examples of SP access from an administration host

The following example shows how to log in to the SP as `naroot`.

```
[admin_host]$ ssh naroot@192.168.123.98
naroot@192.168.123.98's password:
Last login: Thu Jun 3 23:24:37 2010
SP toaster>
```

The following example shows how to log in to the SP with a user account, `joe`, which has been set up on the storage system to have the `login-sp` capability.

```
[admin_host]$ ssh joe@192.168.123.98
joe@192.168.123.98's password:
Last login: Thu Jun 3 23:24:37 2010
SP toaster>
```

Accessing the SP from the system console

You can access the SP from the system console to perform monitoring or troubleshooting tasks.

Step

1. To access the SP CLI from the system console, press Ctrl-G at the storage system prompt.

The SP prompt appears, indicating that you have access to the SP CLI.

Note: To return to the system console, press Ctrl-D and then press Enter.

Example of accessing the SP CLI from the system console

The following example shows the result of pressing Ctrl-G from the system console to access the SP CLI. The `help system power` command is entered at the SP prompt, followed by Ctrl-D and Enter to return to the system console.

```
toaster>
```

(Press Ctrl-G to access the SP CLI.)

```
SP toaster>
```

```
SP toaster> help system power
```

```
system power cycle - power the system off, then on
```

```
system power off - power the system off
```

```
system power on - power the system on
```

```
system power status - print system power status
```

```
SP toaster>
```

(Press Ctrl-D and then Enter to return to the system console.)

```
toaster>
```

SP CLI and system console sessions

Only one administrator can log in to an active SP CLI session at a time. However, the SP allows you to open both an SP CLI session and a separate system console session simultaneously.

The SP prompt appears with SP 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 SP session is `SP toaster>`.

If an SP CLI session is currently open, you or another administrator with privileges to log in to the SP can close the SP CLI session and open a new one. This feature is convenient if you logged in to the SP from one computer and forgot to close the session before moving to another computer, or if another administrator takes over the administration tasks from a different computer.

You can use the SP's `system console` command to connect to the storage system console from the SP. You can then start a separate SSH session for the SP CLI, leaving the system console session active. When you type Ctrl-D to exit from the storage system console, you automatically return to the SP CLI session. If an SP CLI session already exists, the following message appears:

```
User username has an active console 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 SP's system event log.

How to use the SP CLI

The SP CLI provides commands that enable you to remotely access and administer the storage system and diagnose error conditions.

Using the SP CLI, you can perform the following key tasks:

- Remotely administer the storage system by using the Data ONTAP CLI through the SP-redirected system console
- Remotely access the storage system and diagnose error conditions even if the storage system has failed, by performing the following tasks:

- Obtain status information about environmental sensors
- View the storage system console messages, captured in the SP's console log
- View storage system events, captured in the SP'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

Note: There are no man pages for the SP CLI commands.

Next topics

Using online help at the SP CLI on page 198

What you can do in SP admin mode on page 199

What you can do in SP advanced mode on page 202

Connecting to the system console from the SP on page 202

Related concepts

Data ONTAP command-line interface on page 37

Data ONTAP commands at different privilege levels on page 39

Using online help at the SP CLI

The SP online help displays the SP CLI commands and options when you enter the question mark (?) or `help` at the SP prompt.

Steps

1. To display help information for the SP commands, enter one of the following at the SP prompt:

- `help`
- `?`

Example

The following example shows the SP CLI online help:

```
SP toaster> help
date - print date and time
exit - exit from the SP command line interface
events - print system events and event information
help - print command help
priv - show and set user mode
sp - commands to control the SP
rsa - commands for Remote Support Agent
```

```
system - commands to control the system
version - print SP version
```

For more information about the RSA command, see the *Remote Support Agent Concepts and Customer Usage Guide*.

2. To display help information for the option of an SP command, enter the following command at the SP prompt:

```
help SP_command
```

Example

The following example shows the SP CLI online help for the SP events command:

```
SP toaster> help events
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
```

Related concepts

[The Remote Support Agent as a firmware upgrade](#) on page 273

What you can do in SP admin mode

The admin-level SP commands enable you to display system events and logs, reboot the storage system or the SP, create a system core dump, and display status information for system power, system batteries, system sensors, field-replaceable units (FRUs), or the SP.

The following list shows the SP commands that you can enter in admin mode at the SP prompt:

<code>date</code>	Displays system date and time.
<code>events {all info newest <i>number</i> oldest <i>number</i> search <i>keyword</i>}</code>	Displays storage system events that are logged by the SP.
<code>exit</code>	Exits from the SP CLI.
<code>help [<i>command</i>]</code>	Displays a list of available commands. If a command is specified, displays the subcommands available for that command or its syntax.
<code>priv set {admin advanced diag}</code>	Sets the privilege level to access the specified mode for the SP CLI. Attention: You should use advanced or diag commands only under the guidance of technical support.
<code>priv show</code>	Displays the current privilege level for the SP CLI.

<code>rsa</code>	<p>Manages the Remote Support Agent (RSA) if it is installed on your storage system.</p> <p>Note: For information about the RSA, see the <i>Remote Support Agent Concepts and Customer Usage Guide</i>.</p>
<code>sp reboot</code>	Reboots the SP.
<code>sp status[-v -d]</code>	<p>Displays SP status and network configuration information.</p> <p>The <code>-v</code> option displays SP statistics in verbose form.</p> <p>The <code>-d</code> option adds SP debug log to the display.</p> <p>Note: The Data ONTAP <code>sysconfig</code> command displays the status for both the storage system and the SP.</p>
<code>sp update image_URL</code>	<p>Updates the SP firmware by using the image at the specified location.</p> <p>Note: <code>image_URL</code> must not exceed 200 characters.</p>
<code>sp uptime</code>	Displays the current time, the length of time the system has been up, and the average number of jobs in the run queue over the last 1, 5, and 15 minutes.
<code>system battery show</code>	<p>Displays system battery information.</p> <p>Note: This command is available only on the 32xx systems.</p>
<code>system console</code>	<p>Logs in to the system console.</p> <p>Note: You use Ctrl-D to exit from the system console and return to the SP CLI.</p>
<code>system core</code>	<p>Creates a system core dump and resets the storage system. This command has the same effect as pressing the Non-maskable Interrupt (NMI) button on a storage system.</p> <p>Note: The SP stays operational as long as the input power to the storage system is not interrupted.</p>
<code>system fru list</code>	Lists all system FRUs and their IDs .
<code>system fru show fru_id</code>	Displays product information for the specified FRU.

	<p>Note: You can display FRU IDs by using the <code>system fru list</code> command.</p>
<p><code>system log</code></p> <p><code>system power {on off cycle}</code></p>	<p>Displays system console logs.</p> <p>Turns the storage system on or off, or performs a power cycle (turning system power off and then back on.) The standby power stays on to keep the SP running without interruption. During the power cycle, a brief pause occurs before power is turned back on.</p>
	<p>Attention: Using the <code>system power off</code> or the <code>system power cycle</code> command is not a substitute for a graceful shutdown using the Data ONTAP <code>halt</code> command. The <code>system power off</code> and the <code>system power cycle</code> commands might cause an improper shutdown of the storage system (also called a dirty shutdown.)</p>
<p><code>system power status</code></p> <p><code>system reset {primary backup current}</code></p>	<p>Displays the status for the system power supply.</p> <p>Resets the storage system by using the specified firmware image.</p>
	<p>Note: The SP stays operational as long as the input power to the storage system is not interrupted.</p>
<p><code>system sensors</code></p>	<p><code>system sensors</code> has an equivalent command, <code>system sensors show</code>. Both <code>system sensors</code> and <code>system sensors show</code> display the status for the environmental sensors, including their states and current values.</p>
<p><code>system sensors get sensor_name</code></p>	<p>Displays the status and details for the specified sensor.</p> <p>Note: You can obtain <code>sensor_name</code> by using the <code>system sensors</code> or the <code>system sensors show</code> command.</p>
<p><code>version</code></p>	<p>Displays the SP hardware and firmware version information.</p>

Related concepts

[The Remote Support Agent as a firmware upgrade](#) on page 273

What you can do in SP advanced mode

The advanced SP commands enable you to perform more tasks than allowed in admin mode, including displaying the SP command history, SP debug file, SP messages file, and FRU data history, and managing battery firmware and automatic update.

In addition to the SP admin commands, additional SP commands are available in advanced mode:

Attention: You should use advanced commands only under the guidance of technical support.

<code>sp log audit</code>	Displays the SP command history.
<code>sp log debug</code>	Displays the SP debug information.
<code>sp log messages</code>	Displays the SP messages file.
<code>system battery auto_update [status enable disable]</code>	Displays the status of battery firmware automatic update, or enables or disables battery firmware automatic update upon next SP boot.
<code>system battery flash image_URL</code>	Updates battery firmware from the image at the specified location. You use <code>system battery flash</code> if the automatic battery firmware upgrade process has failed for some reason.
<code>system battery verify [image_URL]</code>	Compares the current battery firmware image against a specified firmware image. If <code>image_URL</code> is not specified, the default battery firmware image is used for comparison.
<code>system fru log show</code>	Displays the FRU data history log.

Related tasks

[Setting the privilege level](#) on page 40

Connecting to the system console from the SP

The SP's `system console` command enables you to log in to the storage system from the SP.

Steps

1. Enter the following command at the SP prompt:

```
system console
```

The message “Type Ctrl-D to exit” appears.

2. Press Enter to see the storage system prompt.

The storage system prompt appears.

3. To exit from the system console and return to the SP CLI, press Ctrl-D.

Example of connecting to the system console from the SP

The following example shows the result of entering the `system console` command at the SP prompt. The `vol status` command is entered at the storage system prompt, followed by Ctrl-D, which returns you to the SP prompt.

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

(Press Enter to see the storage system prompt.)

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

(Information about all of the existing volumes is displayed.)

(Press Ctrl-D to exit from the system console and return to the SP CLI.)

```
SP toaster>
```

How to use Data ONTAP to manage the SP

You can manage the SP from the storage system by using the Data ONTAP `sp` commands and by changing the AutoSupport settings that are used by the SP. You can also use Data ONTAP options to control SNMP traps for the SP.

Next topics

[Data ONTAP commands for the SP](#) on page 203

[SP and AutoSupport options](#) on page 204

[SP and SNMP traps](#) on page 205

[Enabling or disabling SNMP traps for Data ONTAP and the SP](#) on page 205

[Disabling SNMP traps for only the SP](#) on page 206

Data ONTAP commands for the SP

Data ONTAP provides `sp` commands that you can use to manage the SP, including setting up the SP, rebooting the SP, displaying the status of the SP, testing the SP, and updating the SP firmware.

The following list shows the admin-level Data ONTAP commands that you can use to manage the SP. The `sp` commands are also described in the `na_sp(1)` man page.

`options sp.setup` Displays whether the SP has been configured.

Note: You configure the SP by using the `setup` or the `sp setup` command. If you use the `setup` command, the `sp setup` command is initiated after `setup` finishes running.

<code>sp help</code>	Displays the Data ONTAP <code>sp</code> commands that you can enter at the storage system prompt.
<code>sp reboot</code>	Reboots the SP and causes the SP to perform a self-test. Any console connection through the SP is lost.
<code>sp setup</code>	Initiates the interactive SP setup script. Note: This command is also available at the boot environment prompt.
<code>sp status</code>	Displays the current status and the network configuration of the SP. Note: This command is also available at the boot environment prompt.
<code>sp test autosupport</code>	Sends a test e-mail to all recipients specified with the <code>autosupport .to</code> option. Note: For this command to work, the <code>autosupport .enable</code> and the <code>autosupport .mailhost</code> options must be configured properly.
<code>sp test snmp</code>	Performs SNMP test on the SP, forcing the SP to send a test SNMP trap to all trap hosts specified in the <code>snmp traphost</code> command. For information about initializing SNMP traps, see the <i>Data ONTAP 7-Mode Network Management Guide</i> .
<code>sp update</code>	Updates the SP firmware. Before using this command, you must use the <code>software install</code> command to install the new SP firmware image. For instructions on how to download and update the SP firmware, see the <i>Data ONTAP 7-Mode Upgrade Guide</i> .

Related concepts

[Ways to configure the SP](#) on page 191

SP and AutoSupport options

The SP extends AutoSupport capabilities by sending alerts and notifications through an AutoSupport message. You can manage AutoSupport event messages and alerts by using the autosupport options.

The SP extends AutoSupport capabilities by sending alerts and “down system” notifications through an AutoSupport message when the storage system goes down, regardless of whether the storage system can send AutoSupport messages. Other than generating these messages on behalf of a system

that is down and attaching additional diagnostic information to AutoSupport messages, the SP has no effect on the storage system's AutoSupport functionality. The AutoSupport configuration settings and message content behavior are inherited from Data ONTAP.

The SP uses the settings of the following Data ONTAP options to send event messages and alerts:

- `autosupport.to`
- `autosupport.mailhost`

You use the `autosupport.content` option to change the amount of information displayed by AutoSupport.

Related concepts

[Contents of AutoSupport event messages](#) on page 186

[AutoSupport options](#) on page 177

SP and SNMP traps

If SNMP is enabled for the SP, the SP generates SNMP traps to configured trap hosts for all "down system" events.

You can enable SNMP traps for both Data ONTAP and the SP. You can also disable the SNMP traps for only the SP and leave the SNMP traps for Data ONTAP enabled.

For information about SNMP traps, see the *Data ONTAP 7-Mode Network Management Guide*.

Enabling or disabling SNMP traps for Data ONTAP and the SP

You can use the `snmp.enable` option to enable or disable SNMP traps for both Data ONTAP and the SP.

The `snmp.enable` option is the master control for SNMP traps for both Data ONTAP and the SP. Consider leaving the `snmp.enable` option to `on` to enable SNMP traps for both Data ONTAP and the SP.

Step

1. To enable or disable SNMP traps for both Data ONTAP and the SP, enter the following command at the storage system prompt:

```
options snmp.enable [on|off]
```

The default is `on`.

Related tasks

[Disabling SNMP traps for only the SP](#) on page 206

Disabling SNMP traps for only the SP

You can disable SNMP traps for only the SP and leave SNMP traps for Data ONTAP enabled.

Step

1. To disable SNMP traps for only the SP, enter the following command at the storage system prompt:

```
options sp.snmp.traps off
```

The default is on.

If the `sp.snmp.traps` option is set to `off`, every time the system boots, an EMS message occurs to inform you that the SNMP trap support for the SP is currently disabled and that you can set the `sp.snmp.traps` option to `on` to enable it. This EMS message also occurs when the `sp.snmp.traps` option is set to `off` and you try to run a Data ONTAP command to use the SP to send an SNMP trap.

You cannot enable SNMP traps for only the SP when SNMP traps for Data ONTAP is disabled. If you set `options snmp.enable` to `off`, both Data ONTAP and the SP stop sending SNMP traps, even if `options sp.snmp.traps` is set to `on`. That is, the following command combination does not result in enabled SNMP traps for only the SP:

```
options snmp.enable off
options sp.snmp.traps on
```

Related tasks

[Enabling or disabling SNMP traps for Data ONTAP and the SP](#) on page 205

How the SP sensors help you monitor system components

There are two types of SP sensors: threshold-based sensors and discrete sensors. Their status information (displayed by the `system sensors` command output) helps you monitor the environmental components of your system.

Note: `system sensors` has an equivalent command, `system sensors show`. Both commands display the same output.

Next topics

[How to determine the status of a threshold-based SP sensor](#) on page 206

[How to determine the status of a discrete SP sensor](#) on page 208

How to determine the status of a threshold-based SP sensor

Threshold-based sensors take periodic readings of a verity of system components. The SP compares the reading of a threshold-based sensor against its preset threshold limits that define a component's

acceptable operating conditions. Based on the sensor reading, the SP displays the sensor state to help you monitor the condition of the component.

Examples of threshold-based sensors include sensors for the system temperatures, voltages, currents, and fan speeds. The specific list of threshold-based sensors depends on the platform.

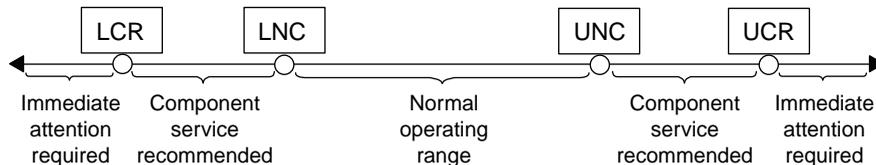
Threshold-based sensors have the following thresholds, displayed in the output of the SP command `system sensors`:

- lower critical (LCR)
- lower noncritical (LNC)
- upper noncritical (UNC)
- upper critical (UCR)

A sensor reading between LNC and LCR or between UNC and UCR means that the component is showing signs of problem and a system failure might occur as a result. Therefore, you should plan for component service soon.

A sensor reading below LCR or above UCR means that the component is malfunctioning and a system failure is about to occur. Therefore, the component requires immediate attention.

The following diagram illustrates the severity ranges that are specified by the thresholds:



You can find the reading of a threshold-based sensor under the Current column in the `system sensors` command output. As the reading of a threshold-based sensor crosses the noncritical and critical threshold ranges, the sensor reports a problem of increasing severity. When the reading exceeds a threshold limit, the sensor's status in the `system sensors` command output changes from `ok` to either `nc` (noncritical) or `cr` (critical), and an event message is logged in the SEL event log.

Some threshold-based sensors do not have all four threshold levels. For those sensors, the missing thresholds show `na` as their limits in the `system sensors` command output. `na` means that the particular sensor has no limit or severity concern for the given threshold, and the SP does not monitor the sensor for that threshold.

Example of the `system sensors` command output

The following example shows the information displayed by the `system sensors` command:

```
SP> system sensors
```

Sensor Name	Current	Unit	Status	LCR	LNC	UNC	UCR
CPU0_Temp_Margin	-55.000	degrees C	ok	na	na	-5.000	0.000
CPU1_Temp_Margin	-56.000	degrees C	ok	na	na	-5.000	0.000
In_Flow_Temp	32.000	degrees C	ok	0.000	10.000	42.000	52.000
Out_Flow_Temp	38.000	degrees C	ok	0.000	10.000	59.000	68.000
PCI_Slot_Temp	40.000	degrees C	ok	0.000	10.000	56.000	65.000
NVMEM_Bat_Temp	32.000	degrees C	ok	0.000	10.000	55.000	64.000

LM56_Temp	38.000	degrees C	ok	na	na	49.000	58.000
CPU0_Error	0x0	discrete	0x0180	na	na	na	na
CPU0_Therm_Trip	0x0	discrete	0x0180	na	na	na	na
CPU0_Hot	0x0	discrete	0x0180	na	na	na	na
CPU1_Error	0x0	discrete	0x0180	na	na	na	na
CPU1_Therm_Trip	0x0	discrete	0x0180	na	na	na	na
CPU1_Hot	0x0	discrete	0x0180	na	na	na	na
IO_Mid1_Temp	30.000	degrees C	ok	0.000	10.000	55.000	64.000
IO_Mid2_Temp	30.000	degrees C	ok	0.000	10.000	55.000	64.000
CPU_VTT	1.106	Volts	ok	1.028	1.048	1.154	1.174
CPU0_VCC	1.154	Volts	ok	0.834	0.844	1.348	1.368
CPU1_VCC	1.086	Volts	ok	0.834	0.844	1.348	1.368
1.0V	0.989	Volts	ok	0.941	0.951	1.057	1.067
1.05V	1.048	Volts	ok	0.980	0.999	1.106	1.125
1.1V	1.096	Volts	ok	1.028	1.038	1.154	1.174
1.2V	1.203	Volts	ok	1.125	1.135	1.261	1.280
1.5V	1.513	Volts	ok	1.436	1.455	1.571	1.591
1.8V	1.754	Volts	ok	1.664	1.703	1.896	1.935
2.5V	2.543	Volts	ok	2.309	2.356	2.621	2.699
3.3V	3.323	Volts	ok	3.053	3.116	3.466	3.546
5V	5.002	Volts	ok	4.368	4.465	5.490	5.636
STBY_1.8V	1.794	Volts	ok	1.678	1.707	1.892	1.911
...							

Example of the `system sensors get sensor_name` command output for a threshold-based sensor

The following example shows the result of entering `system sensors get sensor_name` for the threshold-based sensor 5V:

```
SP> system sensors get 5V
Locating sensor record...
Sensor ID          : 5V (0x13)
Entity ID          : 7.97
Sensor Type (Analog) : Voltage
Sensor Reading     : 5.002 (+/- 0) Volts
Status             : ok
Lower Non-Recoverable : na
Lower Critical      : 4.246
Lower Non-Critical  : 4.490
Upper Non-Critical  : 5.490
Upper Critical      : 5.758
Upper Non-Recoverable : na
Assertion Events    :
Assertions Enabled  : lnc- lcr- ucr+
Deassertions Enabled : lnc- lcr- ucr+
```

How to determine the status of a discrete SP sensor

The Status column of the `system sensors` command output shows the discrete sensors' conditions in hexadecimal values. To interpret the status values of most discrete sensors, you can use the `system sensors get sensor_name` command.

Discrete sensors do not have thresholds. Their readings (displayed under the Current column in the `system sensors` command output) do not carry actual meanings and thus are ignored by the SP.

Examples of discrete sensors include sensors for the fan present, power supply unit (PSU) fault, and system fault. The specific list of discrete sensors depends on the platform.

While the `system sensors get sensor_name` command displays the status information for most discrete sensors, it does not provide status information for the `System_FW_Status`, `System_Watchdog`, `PSU1_Input_Type`, and `PSU2_Input_Type` discrete sensors. However, you can use the following information to interpret these sensors' status values.

System_FW_Status

The System_FW_Status sensor's condition appears in the form of `0xAABB`. You can combine the information of `AA` and `BB` to determine the condition of the sensor.

`AA` can have one of the following values:

- 01** System firmware error
- 02** System firmware hang
- 04** System firmware progress

`BB` can have one of the following values:

- 00** System software has properly shut down
- 01** Memory initialization in progress
- 02** NVMEM initialization in progress (when NVMEM is present)
- 04** Restoring memory controller hub (MCH) values (when NVMEM is present)
- 05** User has entered Setup
- 13** Booting the operating system or LOADER
- 1F** BIOS is starting up
- 20** LOADER is running
- 21** LOADER is programming the primary BIOS firmware. You must not power down the system.
- 22** LOADER is programming the alternate BIOS firmware. You must not power down the system.
- 2F** Data ONTAP is running
- 60** SP has powered off the system
- 61** SP has powered on the system
- 62** SP has reset the system
- 63** SP watchdog power cycle
- 64** SP watchdog cold reset

For instance, the System_FW_Status sensor status `0x042F` means "system firmware progress (04), Data ONTAP is running (2F)."

System_Watchdog

The System_Watchdog sensor can have one of the following conditions:

0x0080	The state of this sensor has not changed
0x0081	Timer interrupt
0x0180	Timer expired
0x0280	Hard reset
0x0480	Power down
0x0880	Power cycle

For instance, the `System_Watchdog` sensor status `0x0880` means a watchdog timeout occurs and causes a system power cycle.

PSU1_Input_Type and PSU2_Input_Type

For direct current (DC) power supplies, the `PSU1_Input_Type` and `PSU2_Input_Type` sensors do not apply. For alternating current (AC) power supplies, the sensors' status can have one of the following values:

0x01xx	220V PSU type
0x02xx	110V PSU type

For instance, the `PSU1_Input_Type` sensor status `0x0280` means that the sensor reports that the PSU type is 110V.

Examples of the `system sensors get sensor_name` command output for discrete sensors

The following examples show the results of entering `system sensors get sensor_name` for the discrete sensors `CPU0_Error` and `IO_Slot1_Present`:

```
SP> system sensors get CPU0_Error
Locating sensor record...
Sensor ID       : CPU0_Error (0x67)
Entity ID      : 7.97
Sensor Type (Discrete): Temperature
States Asserted : Digital State
                  [State Deasserted]

SP> system sensors get IO_Slot1_Present
Locating sensor record...
Sensor ID       : IO_Slot1_Present (0x74)
Entity ID      : 11.97
Sensor Type (Discrete): Add-in Card
States Asserted : Availability State
                  [Device Present]
```

SP commands for troubleshooting the storage system

When you encounter a problem with the storage system, you can use the SP to display information about the problem, create a system core dump, and reboot the storage system, even if the storage system's firmware is corrupted.

The following table describes the common SP commands that you can use at the SP prompt to troubleshoot a storage system.

Condition	Goal	SP command
An environmental sensor has reached an abnormal condition	Display the status for all environmental sensors, their states, and the current values	<code>system sensors show</code>
	Display the status and details for a specific sensor	<code>system sensors get <i>sensor_name</i></code>
The storage system is not responding properly	Access the storage system console from the SP	<code>system console</code>
	Create a system core dump and reboot the system	<code>system core</code>
	Power-cycle the storage system	<code>system power cycle</code>
You receive an AutoSupport message indicating an event such as a hardware component failure or storage system panic	Display what has occurred at the storage system console	<code>system log</code>
	Display all events	<code>events all</code>
	Display a specific number of recent events	<code>events newest <i>number</i></code>
	Search for specific events regarding <i>keyword</i>	<code>events search <i>keyword</i></code>
The storage system firmware is corrupted	Boot the storage system by using the backup image of the storage system firmware	<code>system reset backup</code>
A FRU is malfunctioning	Display the FRU's product information	<code>system fru list</code> to list all FRU IDs <code>system fru show <i>fru_id</i></code> to display product information for a specific FRU

Related references

[What you can do in SP admin mode](#) on page 199

[What you can do in SP advanced mode](#) on page 202

System event log and the SP

The SP has a nonvolatile memory buffer that stores up to 4,000 system events in a system event log (SEL). The SEL stores each audit log entry as an audit event. It is stored in onboard flash memory on the SP.

You can view the audit log entries that are stored in the SEL, along with other stored events, by using the SP `events` command. You can also use the SP `sp log audit` command to perform a quick search for audit log entries in the SEL.

Note: The SEL stores platform-specific events. This log is self-contained and does not support the Syslog Translator.

The primary purpose of the SEL is to help you diagnose system issues. The event list from the SEL is automatically sent by the SP to specified recipients through an AutoSupport message.

The SEL contains the following data:

- Hardware events detected by the SP—for example, system sensor status about power supplies, voltage, or other components
- Errors detected by the SP—for example, a communication error, a fan failure, or a memory or CPU error
- Critical software events sent to the SP by the storage system—for example, a system panic, a communication failure, a boot failure, or a user-triggered “down system” as a result of issuing the SP `system reset` or `system power cycle` command

Note: The SEL uses the SP’s clock to time-stamp events. The SP 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 'pre-init time'.

The following examples show the results of entering the SP `events search keyword` command:

```
SP toaster> events search reboot
Record 3460: Sun Mar 21 07:08:27 2010 [SP CLI.notice]: naroot "sp reboot"
Record 3516: Sun Mar 21 18:11:01 2010 [SP CLI.notice]: root "sp reboot"
Record 3688: Wed Mar 24 04:56:04 2010 [SP CLI.notice]: naroot "sp reboot"
Record 3726: Wed Mar 24 05:12:57 2010 [Trap Event.critical]: SNMP
abnormal_reboot (28)
Record 827: Sun Mar 28 01:07:10 2010 [ONTAP.notice]: Appliance user
command reboot.
...
```

```
SP toaster> events search cycle
Record 3819: Wed Mar 24 23:10:53 2010 [SP CLI.notice]: naroot "system
power cycle "
Record 3820: Wed Mar 24 23:11:02 2010 [IPMI Event.critical]: System power
cycle
Record 3821: Wed Mar 24 23:11:02 2010 [Trap Event.notice]: SNMP
power_cycle_via_sp (24)
Record 3826: Wed Mar 24 23:11:33 2010 [ASUP.notice]: First notification
email |(USER_TRIGGERED (system power cycle)) NOTICE | Sent
```

...

Console log and the SP

The SP monitors the system console regardless of whether administrators are logged in or connected to the console. When system messages are sent to the console, the SP stores them in the console log.

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 SP has power from either of the storage system's power supplies. Since the SP 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 that are sent by the SP.

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

```
SP toaster> system log
Wed Mar 31 18:24:24 GMT [asup.post.host:info]: AutoSupport (HA Group
Notification from partner_node (REBOOT (halt command)) INFO) cannot
connect to url asuppost.company.com (specified host not found)
Wed Mar 31 19:50:30 GMT [ses.shelf.unsupportAllowErr:ALERT]: Unsupported
disk shelf found on channel 0c.
Wed Mar 31 19:50:31 GMT [ses.access.noEnclServ:CRITICAL]: No Enclosure
Services detected through channel 0c.
Wed Mar 31 20:00:00 GMT [monitor.shelf.configError:CRITICAL]: Enclosure
services has detected an error in access to shelves or shelf configuration
0c.
...
```

AutoSupport messages for systems with the SP

For storage systems with the SP, there are two additional types of AutoSupport messages—SP-generated AutoSupport messages about the storage system, and storage system-generated AutoSupport messages about the SP.

SP-generated AutoSupport messages include the following information:

- In the subject line—A notification from the SP of the storage system, listing the system condition or event that caused the AutoSupport message and the log level.
- In the message body—The SP configuration and version information, the storage system ID, serial number, model, and host name.
- In the attachments—The system event logs, the system sensor state as determined by the SP, and the console logs. (The console logs are omitted if the `autosupport.content` option is set to `minimal`.)

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

- The storage system reboots unexpectedly.
- The storage system stops communicating with the SP.
- A watchdog reset occurs.

The watchdog is a built-in hardware sensor that monitors the storage system for a hung or unresponsive condition. If the watchdog detects this condition, it resets the storage system so that the system can automatically reboot and resume functioning. This feature is sometimes called automatic server restart.

When the SP detects a watchdog-generated event on the storage system, it logs this event and, if needed, sends an AutoSupport alert for this event.

- The storage system is power-cycled.
- Firmware power-on self-test (POST) errors occur.
- A user-initiated AutoSupport message occurs.

Storage system-generated AutoSupport messages include the following information:

- In the subject line—A notification from the storage system, including a description of the SP condition or event that caused the AutoSupport message and the log level.
- 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 attachments—Messages from `EMS`, `rc`, `exports`, `hosts`, `resolv_conf`, `nsswitch_conf`, and `cm_stats`

Typical storage system-generated AutoSupport messages about the SP occur under the following conditions:

- The SP stops communicating with the storage system.
- The SP software fails.
- The SP hardware fails.

Related concepts

[Contents of AutoSupport event messages](#) on page 186

[AutoSupport messages](#) on page 185

How to update the SP firmware

You can download and update the SP firmware from the Data ONTAP CLI or the SP CLI.

For instructions on how to download and update the SP firmware, see the *Data ONTAP 7-Mode Upgrade Guide*.

Troubleshooting SP connection problems

If you are having difficulty connecting to the SP, you should verify that your administration host has a secure shell client that supports SSHv2 and that the IP configuration is correct.

Steps

1. Verify that the administration host that you are using to connect to the SP has a secure shell client that supports SSHv2.
2. From the storage system, verify that the SP is online and that the IP configuration is correct by entering the following command at the storage system prompt:

```
sp status
```

3. From the administration host, test the network connection for the SP by entering the following command:

```
ping SP_IP_address
```

4. If the ping fails, do one of the following:
 - Verify that the SP 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 that the SP has a valid IP address. To use the DHCP server or change the IP address for the SP, enter the following command at the storage system prompt:

```
sp setup
```

- Verify that the administration host has a route to the SP.
5. From the storage system prompt, reboot the SP by entering the following command:

```
sp reboot
```

6. If the SP does not reboot, repeat Steps 2 through 5. If the SP still does not reboot, contact technical support for assistance.

Related concepts

Prerequisites for logging in to the SP on page 0

Using the Remote LAN Module for remote system management

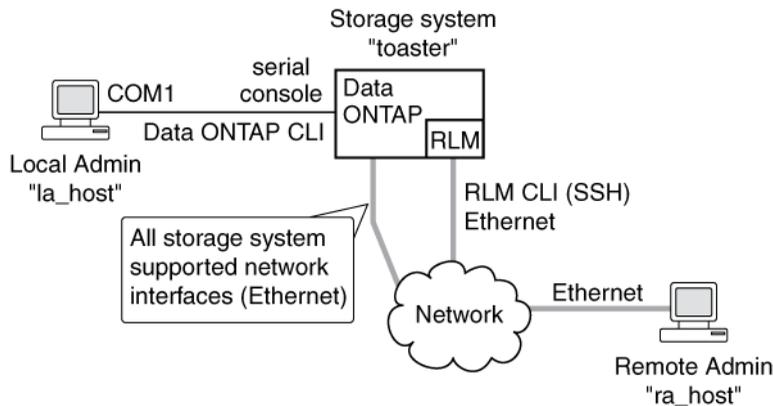
The Remote LAN Module (RLM) is a remote management card that is supported on the 30xx, 31xx, and 60xx storage systems. The RLM provides remote platform management capabilities, including remote access, monitoring, troubleshooting, logging, and alerting features.

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 RLM has a single temperature sensor to detect ambient temperature around the RLM board. Data generated by this sensor is not used for any system or RLM environmental policies. It is only used as a reference point that might help you troubleshoot storage system issues. For example, it might help a remote system administrator determine if a system was shut down due to an extreme temperature change in the system.

For instructions on how to cable your storage system to the RLM, see the *Installing or Replacing a Remote LAN Module* flyer.

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.

The RLM is supported by the Operations Manager. See the Operations Manager Online Help for details.

Next topics

What the RLM does on page 217

Ways to configure the RLM on page 218

How to log in to the RLM on page 221

How to manage the storage system with the RLM on page 225

How to manage the RLM with Data ONTAP on page 231

How to display information about the storage system and the RLM on page 234

Comparison of Data ONTAP and RLM commands on page 239

How to troubleshoot the storage system with the RLM on page 241

How to update the RLM firmware on page 241

How to troubleshoot RLM problems on page 242

Related concepts

The e0M interface on page 45

Related information

The NOW site - <http://now.netapp.com/>

What the RLM does

The RLM command line interface (CLI) commands enable you to remotely access and administer the storage system and diagnose error conditions. Also, the RLM extends AutoSupport capabilities by sending alerts and notifications through an AutoSupport message.

Using the RLM CLI commands, you can perform the following tasks:

- Remotely administer the storage system by using the Data ONTAP CLI through 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 the 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

The RLM extends AutoSupport capabilities by sending alerts and “down system” or “down filer” notifications through an AutoSupport message when the storage system goes down, regardless of

whether the storage system can send AutoSupport messages. Other than generating these messages on behalf of a system that is down, and attaching additional diagnostic information to AutoSupport messages, the RLM has no effect on the storage system's AutoSupport functionality. The AutoSupport configuration settings and message content behavior of the RLM are inherited from Data ONTAP.

In addition to AutoSupport messages, the RLM generates SNMP traps to configured trap hosts for all “down system” or “down filer” events, if SNMP is enabled for the RLM.

Hardware-assisted takeover is available on systems that support the RLM and have the RLM modules set up. For more information about hardware-assisted takeover, see the *Data ONTAP 7-Mode High-Availability Configuration Guide*.

The RLM supports the SSH protocol for CLI access from UNIX clients and PuTTY for CLI access from PC clients. Telnet and RSH are not supported by the RLM, and system options to enable or disable them have no effect on the RLM.

Note: The RLM ignores the `autologout.telnet.timeout` and the `autologout.console.timeout` options. The settings for these options do not have any effect on the RLM.

Related concepts

[How to troubleshoot the storage system with the RLM](#) on page 241

[The AutoSupport feature](#) on page 175

Ways to configure the RLM

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

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

- Initializing a storage system that has the RLM pre-installed
When the storage system setup process is complete, the `rlm setup` command runs automatically. For more information about the entire setup process, see the *Data ONTAP 7-Mode Software Setup Guide*.
- Running the Data ONTAP `setup` script
The `setup` script ends by initiating the `rlm setup` command.
- Running the Data ONTAP `rlm setup` command

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

Next topics

[Prerequisites for configuring the RLM](#) on page 219

[Configuring the RLM](#) on page 219

Prerequisites for configuring the RLM

Before you configure the RLM, you must gather information about your network and your AutoSupport settings.

The following is the information you need to gather:

- Network information
You can configure the RLM using DHCP or static addressing. If you are using a static IP address for the RLM, you need the following information:

- An available static IP address
- The netmask of your network
- The gateway of your network

- AutoSupport information

The RLM sends event notifications based on the following AutoSupport settings:

- `autosupport.to`
- `autosupport.mailhost`

It is best that you set at least the `autosupport.to` option before configuring the RLM. You are prompted to enter the name or the IP address of the AutoSupport mail host when you configure the RLM.

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

Related tasks

[Configuring AutoSupport](#) on page 177

Configuring the RLM

You can use the `setup` command or the `rlm setup` command to configure the RLM.

It is best to configure AutoSupport before configuring the RLM. Data ONTAP automatically sends AutoSupport configuration to the RLM, allowing the RLM to send alerts and notifications through an AutoSupport message.

You must use either a static or a DHCP address for RLM configuration. If no static or DHCP address is configured, the RLM has no network connectivity.

Steps

1. At the storage system prompt, enter one of the following commands:

- `setup`
- `rlm setup`

If you enter `setup`, the `rlm setup` script starts automatically after the `setup` command runs.

2. When the RLM setup asks you whether to configure the RLM, enter `y`.
3. Do one of the following when the RLM setup asks you whether to enable DHCP on the RLM.
 - To use DHCP addressing, enter `y`.
 - To use static addressing, enter `n`.
4. If you do not enable DHCP for the RLM, the RLM setup prompts you for static IP information. Provide the following information when prompted:
 - The IP address for the RLM
 - The netmask for the RLM
 - The IP address for the RLM gateway
 - The name or IP address of the mail host to use for AutoSupport (if you use the `setup` command.)
5. At the storage system prompt, enter the following command to verify that the RLM network configuration is correct:

```
rlm status
```

6. At the storage system prompt, enter the following command to verify that the RLM AutoSupport function is working properly:

```
rlm test autosupport
```

Note: The RLM uses the same mail host information that Data ONTAP uses for AutoSupport. The `rlm test autosupport` command requires that you set up the `autosupport.to` option properly.

The following message is a sample of the output Data ONTAP displays:

```
Sending email messages via SMTP server at mailhost@companyname.com. If
autosupport.enable is on, then each email address in autosupport.to
should receive the test message shortly.
```

Examples for configuring the RLM and displaying the configuration information

The following example shows that the RLM is configured to use a static IP address:

```
storage-system> rlm setup
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
Verifying mailhost settings for RLM use...
```

The following example shows that the RLM is configured to use DHCP:

```
storage-system> rlm setup
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
Verifying mailhost settings for RLM use...
```

The following example displays the RLM status and configuration information:

```
storage-system> rlm status
Remote LAN Module           Status: Online
  Part Number:              110-00030
  Revision:                 A0
  Serial Number:            123456
  Firmware Version:        3.0
  Mgmt MAC Address:         00:A0:98:01:7D:5B
  Ethernet Link:            up, 100Mb, full duplex, auto-neg complete
  Using DHCP:               no
  IP Address:               192.168.123.98
  Netmask:                  255.255.255.0
  Gateway:                  192.168.123.1
```

Related concepts

[Prerequisites for configuring the RLM](#) on page 219

[The AutoSupport feature](#) on page 175

How to log in to the RLM

To log in to the RLM, you must install a Secure Shell client application and ensure that you have administrative privileges on the storage system.

The following are the prerequisites for logging in to the RLM:

- A Secure Shell client application that is appropriate for your administration host, such as SSH, OpenSSH for UNIX hosts, or PuTTY for Windows hosts
The RLM accepts only SSH connections. It does not respond to other protocols.
- The RLM's naroot account or a Data ONTAP user account with the credentials of the admin role or a role with the login-sp capability

If the RLM is running firmware version 4.0 or later and is configured to use an IPv4 address, the RLM rejects SSH login requests and suspends all communication with the IP address for 15 minutes if five SSH login attempts fail repeatedly within 10 minutes. The communication resumes after 15 minutes, and you can try to log in to the RLM again.

Next topics

Accounts that can access the RLM on page 222

Logging in to the RLM from a UNIX host on page 223

Logging in to the RLM from a Windows host on page 224

RLM CLI and system console sessions on page 225

Related concepts

How to manage administrator and diagnostic access on page 115

Secure protocols and storage system access on page 49

Accounts that can access the RLM

The RLM comes with an account named `naroot`. Only the RLM's `naroot` account and Data ONTAP user accounts with the credentials of the `admin` role or a role with the `login-sp` capability can log in to the RLM. These users have access to all commands available on the RLM.

For enhanced security, the RLM does not allow you to log in with the Data ONTAP account name `root`. Instead, it maps the Data ONTAP `root` account to the RLM `naroot` account. You use the RLM's `naroot` account and the Data ONTAP `root` password to log into the RLM.

Note: If you disable the `root` account's access to the storage system, the RLM's `naroot` access to the storage system is automatically disabled.

You cannot create user accounts directly from the RLM. However, users created in Data ONTAP with the credentials of the `admin` role or a role with the `login-sp` capability can log in to the RLM. Changes to user account credentials on the storage system are automatically updated to the RLM.

You cannot use the following generic names as account names to access the RLM. Therefore, it is best not to use them as Data ONTAP account names or assign them to Data ONTAP groups that have the `admin` role or a role that includes the `login-sp` capability.

- `adm`
- `bin`
- `cli`
- `daemon`
- `ftp`
- `games`
- `halt`
- `lp`
- `mail`
- `man`
- `netapp`
- `news`
- `nobody`
- `operator`

- shutdown
- sshd
- sync
- sys
- uucp
- www

Related concepts

How to manage administrator and diagnostic access on page 115

Predefined roles on page 128

Supported capability types on page 129

Related tasks

Creating a new role and assigning capabilities to roles on page 131

Modifying an existing role or its capabilities on page 132

Disabling root access to the storage system on page 119

Logging in to the RLM from a UNIX host

You can log in to the RLM from a UNIX host, if a Secure Shell application is installed on the UNIX host.

Ensure that a secure shell application is installed on the UNIX host.

RLM firmware version 4.0 or later accepts only SSHv2 access to the RLM. You must ensure that the UNIX host you use to access the RLM supports SSHv2.

Step

1. Enter the following command from the UNIX host:

```
ssh username@RLM_IP_address
```

Examples of RLM access from a UNIX host

The following example shows how to log in to the RLM as naroot.

```
ssh naroot@192.168.123.98
```

The following example shows how to log in to the RLM with a user account, joe, which has been set up on the storage system.

```
ssh joe@192.168.123.98
```

Logging in to the RLM from a Windows host

You can log in to the RLM from a Windows host, if a Secure Shell application for Windows, such as PuTTY, is installed.

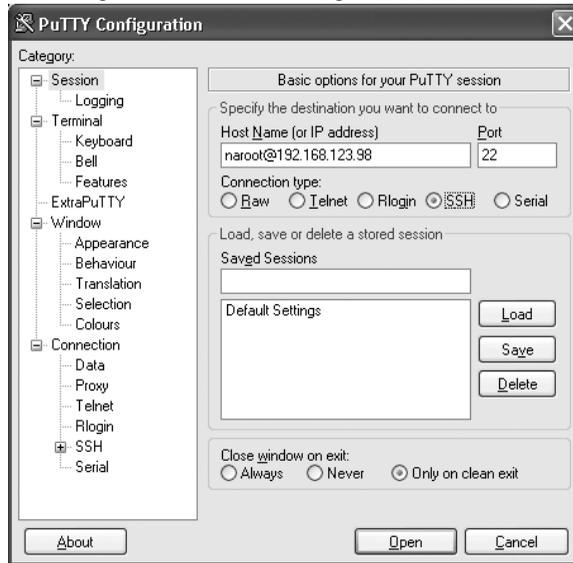
RLM firmware version 4.0 or later accepts only SSHv2 access to the RLM. You must ensure that the Windows host you use to access the RLM supports SSHv2.

Steps

1. Start a Windows session for the Secure Shell application.

Example

You can use the PuTTY Configuration window to log in to the RLM:



2. Enter the IP address of the RLM.
3. Ensure that the SSH protocol is selected.
4. Click **Open**.
5. When you are prompted, use the RLM `naroot` account or a Data ONTAP user account with the `admin` role credentials to log in to the RLM.

RLM CLI and system console sessions

Only one administrator can log in to an active RLM CLI session at a time. However, the RLM allows you to open both an RLM CLI session and a separate, RLM-redirected system console session simultaneously.

The RLM prompt appears with RLM in front of the host name 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 is `RLM toaster>`.

If an RLM CLI session is currently open, you or another administrator with privileges to log in to the RLM can close the RLM CLI session and open a new one. This feature is convenient if you logged in to the RLM from one computer and forgot to close the session before moving to another computer, or if another administrator takes over the administration tasks from a different computer.

When you use the RLM's `system console` command to connect to the storage system console from the RLM, you can start a separate SSH session for the RLM CLI, leaving the system console session active. When you type `Ctrl-D` to exit from the storage system console, you automatically return to the RLM CLI session. If an RLM CLI session already exists, the following message appears:

```
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.

How to manage the storage system with the RLM

The RLM enables you to manage the storage system by using the RLM CLI. The RLM CLI has the same features available in the Data ONTAP CLI.

The CLI features include:

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

Like the Data ONTAP CLI, the RLM CLI provides two privilege levels, `admin` and `advanced`, with different command sets.

Note: The RLM CLI commands are *not* documented in online command line manual (`man`) pages.

Next topics

[Using online help at the RLM CLI](#) on page 226

[What you can do in RLM admin mode](#) on page 227

[RLM admin mode command syntax summary](#) on page 227

[What you can display in RLM advanced mode](#) on page 229

[Connecting to the storage system console from the RLM](#) on page 229

[Controlling storage system power from the RLM](#) on page 230

Related concepts

[Data ONTAP command-line interface](#) on page 37

[Data ONTAP commands at different privilege levels](#) on page 39

Using online help at the RLM CLI

The RLM online help displays all RLM commands and options when you enter the question mark (?) or `help` at the RLM prompt.

Steps

1. To display help information for RLM commands, enter one of the following at the RLM prompt:

- `help`
- `?`

Example

The following example shows the RLM CLI online help:

```
RLM toaster> help
date - print date and time
exit - exit from the RLM command line interface
events - print system events and event information
help - print command help
priv - show and set user mode
rlm - commands to control the RLM
rsa - commands for Remote Support Agent
system - commands to control the system
version - print RLM version
```

For more information about the RSA command, see the *Remote Support Agent Concepts and Customer Usage Guide*.

2. To display help information for the option of an RLM command, enter the following command at the RLM prompt:

```
help RLM_command
```

Example

The following example shows the RLM CLI online help for the RLM `events` command:

```
RLM toaster> help events
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
```


<code>events {all info newest oldest search string}</code>	Displays storage system events logged by the RLM.
<code>exit</code>	Exits from the RLM command-line interface.
<code>help [command]</code>	Displays a list of available commands. If a command is specified, displays the subcommands available for that command or its syntax.
<code>priv set {admin advanced diag}</code>	Sets the privilege level to access the specified mode.
<code>priv show</code>	Displays the current privilege level.
<code>rlm reboot</code>	Reboots the RLM. This action takes approximately one minute.
<code>rlm sensors [-c]</code>	Displays the RLM environmental sensor status. The <code>-c</code> option, which takes a few seconds to display, shows current values, rather than cached values.
<code>rlm status[-v -d]</code>	Displays RLM status. The <code>-v</code> option displays verbose statistics. The <code>-d</code> option displays RLM debug information. Note: The Data ONTAP <code>sysconfig</code> command displays both the storage system and RLM status.
<code>rlm update http://path [-f]</code>	Updates the RLM firmware. The <code>-f</code> option issues a full image update.
<code>rsa</code>	Manages the RSA if it is installed on your storage system. Note: For information about the RSA, see the <i>Remote Support Agent Concepts and Customer Usage Guide</i> .
<code>system console</code>	Logs in to the Data ONTAP CLI. Use Ctrl-D to exit.
<code>system core</code>	Dumps the storage system core and resets the storage system. This command has the same effect as 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.
<code>system power {on off cycle}</code>	Turns the storage system on or off, or performs a power cycle. Standby power stays on. Using the <code>system power</code> command might cause an improper shutdown of the storage system. During power-cycling, a brief pause occurs before power is turned back on.

<code>system power status</code>	Displays status for each power supply, such as presence, input power, and output power.
<code>system reset {primary backup current}</code>	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.
<code>version</code>	Displays the RLM version information, including hardware and firmware information.

Related concepts

[The Remote Support Agent as a firmware upgrade](#) on page 273

What you can display in RLM advanced mode

The RLM advanced commands display more information than is available in administrative mode, including the RLM command history, the RLM debug file, a list of environmental sensors, and RLM statistics.

The following list shows the RLM commands you can use in advanced mode:

<code>rlm log audit</code>	Displays the RLM command history.
<code>rlm log debug</code>	Displays the RLM debug file.
<code>rlm log messages</code>	Displays the RLM messages file.
<code>system sensors</code>	Displays a list of environmental sensors, their states, and their current values.
<code>rlm status -v</code>	Displays RLM statistics.

Related tasks

[Setting the privilege level](#) on page 40

Connecting to the storage system console from the RLM

The RLM's `system console` command enables you to log in to the storage system from the RLM.

Steps

1. Enter the following command at the RLM prompt:

```
system console
```

The message “Type Ctrl-D to exit” appears.

2. Press Enter to see the storage system prompt.

You use Ctrl-D to exit from the storage system console and return to the RLM CLI.

The storage system prompt appears, and you can enter Data ONTAP commands.

Example of connecting to the storage system console from the RLM

The following example shows the result of entering the `system console` command at the RLM prompt. The `vol status` command is entered at the storage system prompt, followed by Ctrl-D, which returns you to the RLM prompt.

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

(Press Enter to see the storage system prompt.)

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

(Information about all of the existing volumes is displayed.)

(Press Ctrl-D to exit from the storage system console and return to the RLM CLI.)

```
RLM toaster>
```

Controlling storage system power from the RLM

The RLM's `system power` command enables you to turn the power on or off or to power-cycle the storage system remotely.

The `system power cycle` command turns system power off and then back on. The power supplies provide a standby voltage that is always present, even when the storage system is off. This keeps the RLM running without interruption. However, turning the power off or power-cycling the storage system may cause an improper shutdown of the storage system (also called a dirty shutdown).

Steps

1. Enter the following command at the RLM prompt:

```
system power cycle
```

Example

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

A warning message indicates that issuing the `system power` command is not a substitute for a graceful shutdown using the Data ONTAP `halt` command.

2. To turn off the storage system, enter `y`.

A few seconds later, the storage system is turned back on, and the boot environment prompt appears. In the meantime, the RLM prompt awaits your next command.

How to manage the RLM with Data ONTAP

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

If SNMP is enabled, the RLM also generates SNMP traps to configured trap hosts for all “down system” or “down filer” events.

Next topics

[Data ONTAP commands for the RLM](#) on page 231

[RLM and AutoSupport options](#) on page 232

[RLM and SNMP options](#) on page 232

[Enabling or disabling SNMP traps for Data ONTAP and the RLM](#) on page 232

[Disabling SNMP traps for only the RLM](#) on page 233

Data ONTAP commands for the RLM

Data ONTAP provides `rlm` commands that allow you to manage the RLM, including setting up the RLM, rebooting the RLM, displaying the status of the RLM, and updating the RLM firmware.

The following table describes the Data ONTAP commands for the RLM. These commands are also described in the `na_rlm(1)` 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.

Data ONTAP Command for the RLM	Description
<code>rlm help</code>	Displays the list of <code>rlm</code> commands available with the current release of Data ONTAP.
<code>rlm reboot</code>	Reboots the RLM and causes the RLM to perform a self-test. Any console connection through the RLM is lost.
<code>rlm setup</code>	Initiates the interactive RLM setup script.
<code>rlm status</code>	Displays the current status of the RLM.
<code>rlm test autosupport</code>	Sends a test e-mail to all recipients specified with the <code>autosupport .to</code> option.
<code>rlm test snmp</code>	Performs SNMP test on the RLM, forcing the RLM to send a test SNMP trap to all trap hosts specified in the <code>snmp traphost</code> command. For information on initializing SNMP traps, see the <i>Data ONTAP 7-Mode Network Management Guide</i> .

Data ONTAP Command for the RLM	Description
rlm update	Updates the RLM firmware. For instructions on how to download and update the RLM firmware, see the <i>Data ONTAP 7-Mode Upgrade Guide</i> .

Related concepts

[Ways to configure the RLM](#) on page 218

RLM and AutoSupport options

The RLM extends AutoSupport capabilities by sending alerts and notifications through an AutoSupport message. You can manage AutoSupport event messages and alerts by using the autosupport options.

The RLM uses the settings of the following options to send event messages and alerts:

- `autosupport.to`
- `autosupport.mailhost`

You use the `autosupport.content` option to change the amount of information displayed by Data ONTAP and RLM AutoSupport commands.

Related concepts

[Contents of AutoSupport event messages](#) on page 186

[AutoSupport options](#) on page 177

RLM and SNMP options

If SNMP is enabled for the RLM, the RLM generates SNMP traps to configured trap hosts for all "down system" or "down filer" events.

You can enable SNMP traps for both Data ONTAP and the RLM. You can also disable the SNMP traps for only the RLM and leave the SNMP traps for Data ONTAP enabled.

Enabling or disabling SNMP traps for Data ONTAP and the RLM

You can enable or disable SNMP traps for both Data ONTAP and the RLM by using the `snmp.enable` option.

The `snmp.enable` option is the master control for enabling or disabling SNMP traps for both Data ONTAP and the RLM. Consider leaving the `snmp.enable` option set to `on` to enable SNMP traps for both Data ONTAP and the RLM.

Step

1. Enter the following command to enable or disable SNMP traps for both Data ONTAP and the RLM:

```
options snmp.enable [on|off]
```

The default option is on.

Note: If you enable SNMP traps on the storage system and the currently installed RLM firmware version does not support SNMP, an EMS message is logged requesting an upgrade of the RLM firmware. Until the firmware upgrade is performed, SNMP traps are not supported on the RLM. For instructions on how to download and update the RLM firmware, see the *Data ONTAP 7-Mode Upgrade Guide*.

Related tasks

[Disabling SNMP traps for only the RLM](#) on page 233

Disabling SNMP traps for only the RLM

You can disable SNMP traps for only the RLM and leave SNMP traps for Data ONTAP enabled.

Step

1. To disable SNMP traps for only the RLM, enter the following command:

```
options rlm.snmp.traps off
```

The default option is on.

If the `rlm.snmp.traps` option is set to `off`, every time the system boots, an EMS message occurs to inform you that the SNMP trap support for the RLM is currently disabled and that you can set the `rlm.snmp.traps` option to `on` to enable it. This EMS message also occurs when the `rlm.snmp.traps` option is set to `off` and you try to run a Data ONTAP command to use the RLM to send an SNMP trap.

You cannot enable SNMP traps for only the RLM when SNMP traps for Data ONTAP is disabled. If you set `options snmp.enable` to `off`, both Data ONTAP and the RLM stop sending SNMP traps, even if `options rlm.snmp.traps` is set to `on`. That is, the following command combination does not result in enabled SNMP traps for only the RLM:

```
options snmp.enable off
```

```
options rlm.snmp.traps on
```

Related tasks

[Enabling or disabling SNMP traps for Data ONTAP and the RLM](#) on page 232

How to display information about the storage system and the RLM

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 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 AutoSupport messages generated by the RLM. Most of the information is stored in the SEL or in captured console messages.

All log entries are recorded with Coordinated Universal Time (UTC) for the time format.

Note: The RLM does not use the time zone setting from the storage system.

Next topics

[RLM CLI commands that display real-time information](#) on page 234

[How to use the RLM to monitor the storage system during a power cycle](#) on page 236

[System event log and the RLM](#) on page 236

[Console log and the RLM](#) on page 237

[AutoSupport messages for systems with the RLM](#) on page 238

RLM CLI commands that display real-time information

Using the RLM CLI commands in admin mode, you can view the status of the storage system power, the status of the RLM, and the version of the RLM. Using the RLM CLI commands in advanced mode, you can view internal RLM statistics and the RLM environmental sensor.

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

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

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

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

RLM CLI in admin mode

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

The following examples show how information is displayed using commands at the RLM advanced 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
total 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

```

How to use the RLM to monitor the storage system during a power cycle

When you power-cycle the storage system, no real-time messages regarding the boot progress appear in the RLM console. To monitor the storage system during a power cycle, use SSH to log in to the RLM CLI and start a system console session with Data ONTAP. Leave this system console session active and start a second SSH session with the RLM CLI. You can then simultaneously interact with the RLM CLI and access the storage system with the system console.

When you power-cycle the storage system using the RLM, the active session to the system console provides real-time output from the system, including the progress of the system boot.

System event log and the RLM

The RLM has a nonvolatile memory buffer that stores up to 4,000 system events in a system event log (SEL). The SEL stores each audit log entry as an audit event. It is stored in onboard flash memory on the RLM. When the buffer is full, the oldest records are overwritten by the newest records.

You can view the audit log entries that are stored in the SEL, along with other stored events, by entering the RLM `events` command. You can also use the `rlm log audit` command to perform a

quick search for audit logs from the SEL. However, the debug logs and message logs are stored separately on the RLM in its RAM and provide debug data for RLM firmware.

The SEL stores platform-specific events. This log is self-contained and does not support the Syslog Translator.

The primary purpose of the SEL is to help you diagnose system issues. The event list from the SEL is automatically sent by the RLM to specified recipients in an AutoSupport message.

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 boot environment prompt, a 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.

The following example shows the result of entering the RLM `events` command:

```
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 and the RLM

The RLM monitors the storage system console regardless of whether 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 RLM CLI `system log` command.

AutoSupport messages for systems with the RLM

For storage systems with the RLM, there are two additional types of AutoSupport messages—RLM-generated AutoSupport messages about the storage system, and storage system-generated AutoSupport messages about the RLM.

RLM-generated AutoSupport messages include the following information:

- In the subject line—A notification from the RLM of the storage system, listing the system condition or event that caused the AutoSupport message and the log level.
- In the message body—The RLM configuration and version information, the storage system ID, serial number, model, and host name.
- In the attachments—The system event logs, the system sensor state as determined by the RLM, and the console logs. (The console logs are omitted if the `autosupport.content` option is set 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 watchdog is a built-in hardware sensor that monitors the storage system for a hung or unresponsive condition. If the watchdog detects this condition, it resets the storage system so that the system can automatically reboot and resume functioning. This feature is sometimes called automatic server restart.

When the RLM detects a watchdog-generated event occurs on the storage system, it logs this event and, if needed, sends an AutoSupport alert for this event.

- The storage system is power-cycled.
- Firmware power-on self-test (POST) errors occur.
- A user-initiated AutoSupport message occurs.

Storage system-generated AutoSupport messages include the following information:

- In the subject line—A notification from the storage system with the RLM, including a description of the RLM condition or event that caused the AutoSupport message and the log level.
- 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 attachments—Messages from `EMS`, `rc`, `exports`, `hosts`, `resolv_conf`, `nsswitch_conf`, and `cm_stats`

Typical storage system-generated AutoSupport messages about the RLM occur under the following conditions:

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

Related concepts

[Contents of AutoSupport event messages](#) on page 186

[AutoSupport messages](#) on page 185

Comparison of Data ONTAP and RLM commands

Whether you use a Data ONTAP command or an RLM command to manage the RLM depends on the task you want to perform.

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 the RLM in a new storage system	When you power on a storage system for the first time, the <code>setup</code> command begins to run automatically. When the storage system setup process is complete, the <code>rlm setup</code> command follows automatically and prompts you for RLM configuration information.	
Reconfigure the RLM in an existing storage system	<code>setup</code> or <code>rlm setup</code> Note: If you use the <code>setup</code> command, the <code>rlm setup</code> command is initiated after <code>setup</code> finishes running.	
Test the RLM's AutoSupport setting	<code>rlm test autosupport</code>	
Perform SNMP test on the RLM	<code>rlm test snmp</code>	
Display Data ONTAP <code>rlm</code> commands	<code>rlm help</code>	
Log in to the RLM		From a UNIX host, enter, <code>ssh user@RLM_IP_addr</code>
Display RLM CLI commands		<code>help</code> or <code>?</code>
Display the twenty most recent events logged by RLM		<code>events newest 20</code>

Action	Data ONTAP Command or Procedure	RLM Command
Display a summary of information about the records in the events log		<code>events info</code>
Display whether the RLM has been configured	<code>options rlm.setup</code> Note: The RLM is configured through the <code>setup</code> or the <code>rlm setup</code> command.	
Display the RLM configuration	<code>rlm status</code> or <code>sysconfig -v</code> Note: <code>sysconfig -v</code> requires advanced mode.	<code>rlm status</code>
Display statistics gathered by RLM	<code>rlm status -v</code> Note: Requires advanced mode.	<code>rlm status -v</code>
Display the system hardware sensor list		<code>system sensors</code> Note: Requires advanced mode.
Log in to the system to manage storage system resources		<code>system console</code> Note: Use Ctrl-D to exit to the RLM CLI.
Dump the system core and reset the storage system		<code>system core</code>
Reset the RLM	<code>rlm reboot</code>	<code>rlm reboot</code>
Update the RLM firmware	<code>software install http://path/RLM_FW.zip -f</code> <code>rlm update[-f]</code> Note: The <code>-f</code> option of the <code>rlm update</code> command requires advanced mode. For information about when to use <code>-f</code> , see the <i>Data ONTAP 7-Mode Upgrade Guide</i> .	<code>rlm update http://path_hostname/RLM_FW.tar.gz [-f]</code> Note: The <code>-f</code> option issues a full image update.

How to troubleshoot the storage system with the RLM

When you encounter a problem with the storage system, you can use the RLM to display information about the problem, create a system core dump, and reboot the storage system, even if the storage system's firmware is corrupted.

The following table describes the common RLM commands that you can use to troubleshoot a storage system.

Note: If you configure the AutoSupport feature, the RLM sends you status messages about both the storage system and the RLM.

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	<code>system console</code>
You receive an AutoSupport message indicating an event such as a hardware component failure or storage system panic.	Display what has occurred at the storage system console	<code>system log</code>
	Display all events	<code>events all</code>
	Display a specific number of recent events	<code>events newest <i>number</i></code>
	Search for specific events in the SEL	<code>events search <i>string</i></code>
The storage system is hanging	Create a system core dump and reboot the storage system	<code>system core</code>
	Power-cycle the storage system	<code>system power cycle</code>
The storage system firmware is corrupted	Boot the storage system by using a backup copy of the storage system firmware	<code>system reset backup</code>

How to update the RLM firmware

You can download and update the RLM firmware from the Data ONTAP CLI or the RLM CLI.

For instructions on how to download and update the RLM firmware and how to troubleshoot RLM firmware update problems, see the *Data ONTAP 7-Mode Upgrade Guide*.

How to troubleshoot RLM problems

RLM problems might result from communication problems, configuration problems, connection problems, RLM hardware failures, or RLM firmware update problems.

Next topics

[Troubleshooting RLM communication problems](#) on page 242

[Troubleshooting RLM configuration problems](#) on page 242

[Troubleshooting RLM connection problems](#) on page 242

[Troubleshooting RLM hardware failures](#) on page 243

Troubleshooting RLM communication problems

A communication failure between the storage system and the RLM might result in RLM problems.

Step

1. 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)`

Troubleshooting RLM configuration problems

If you are having difficulty configuring the RLM, you should verify that the IP configuration is correct.

Steps

1. Verify the RLM is online and the IP configuration is correct by entering the following command at the storage system prompt:
2. If the RLM is configured using DHCP, reconfigure the RLM using a static IP address by entering the following command at the storage system prompt:

```
rlm status
```

```
rlm setup
```

Troubleshooting RLM connection problems

If you are having difficulty connecting to the RLM, you should verify that you are using a secure shell client and that the IP configuration is correct.

Steps

1. Verify that you are using a secure shell client to connect to the RLM.

- From the storage system, verify the RLM is online and the IP configuration is correct by entering the following command at the storage system prompt:

```
rlm status
```

- From the administration host, test the network connection for the RLM by entering the following command:

```
ping rlm_IP_address
```

- If the ping fails, do one of the following:

- Verify that 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 that the RLM has a valid IP address. At the storage system prompt, enter the `rlm setup` command to use the DHCP server or assign a valid IP address.
- Verify that the administration host has a route to the RLM.

- From the storage system prompt, reboot the RLM by entering the following command:

```
rlm reboot
```

Note: It takes approximately one minute for the RLM to reboot.

- If the RLM does not reboot, repeat Steps 2 through 5. If the RLM still does not reboot, contact technical support for assistance.

Related concepts

[How to log in to the RLM](#) on page 221

Troubleshooting RLM hardware failures

An RLM problem can occur when a hardware failure has occurred on the RLM.

When the RLM fails, an EMS event similar to the following can be found:

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

Steps

- Run diagnostics by entering the following command from the boot environment prompt:

```
boot_diags
```

The diagnostics main menu appears.

all	Run all system diagnostics
mb	motherboard diagnostic
mem	main memory diagnostic
agent	agent & rlm diagnostic
cf-card	CompactFlash controller diagnostic
fcsl	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
           prompt)

```

- From the main menu, enter the following option:

agent

Example

```
Enter Diag, Command or Option: agent
```

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]:

```

- From the RLM diagnostic prompt, enter test number 1.

Example

```
Select test or feature by number [0]: 1
```

Note: It takes approximately ten minutes to complete this test.

This step initiates a comprehensive test that includes running tests 2 through 8 and 11. The results of each test are displayed.

- Based on the results of Step 3, diagnose the problem. If the problem persists, reset the RLM and repeat Steps 1 to 4.

If the problem still persists, replace the RLM.

Using the Baseboard Management Controller for remote system management

The Baseboard Management Controller (BMC) is a remote management device that is built into the motherboard of the FAS20xx storage systems. It provides remote platform management capabilities, including remote access, monitoring, troubleshooting, logging, and alerting features.

The BMC firmware supports Intelligent Platform Management Interface (IPMI) version 2.0, which by default supports Serial Over LAN (SOL) for console redirection.

The BMC stays operational regardless of the operating state of the storage system. Both the BMC and its dedicated Ethernet NIC use a standby voltage for high availability. The BMC is available as long as the storage system has input power to at least one of the storage system's power supplies.

The BMC monitors environmental sensors, including sensors for the temperature of the system's nonvolatile memory (NVMEM) battery, motherboard, and CPU, and for the system's voltage level. When the BMC detects that an environmental sensor has reached a critically low or critically high state, it generates AutoSupport messages and shuts down the storage system. The data generated by the sensors can be used as a reference point to help you troubleshoot storage system issues. For example, it can help a remote system administrator determine if a system was shut down due to an extreme temperature change in the system.

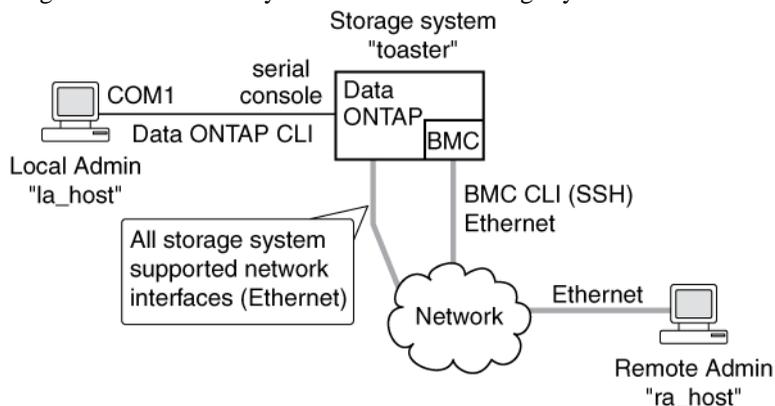
The BMC also monitors non-environmental sensors for the status of the BIOS, power, CPU, and serial-attached SCSI (SAS) disks. These sensors are recorded by the BMC to assist support personnel.

The following table lists the names and the description of the sensors that BMC monitors:

Sensor Name	Description
1.1V	Board 1.1V sensor
1.2V	Board 1.2V sensor
1.5V	Board 1.5V sensor
1.8V	Board 1.8V sensor
2.5V	Board 2.5V sensor
3.3V	Board 3.3V sensor
CPU 1.2V	CPU 1.2V sensor
12.0V	Power 12V sensor
BIOS Status	BIOS status normal
Board Temp Top	Temperature at the top side of the board
Board Temp Bot	Temperature at the bottom side of the board

Sensor Name	Description
CPU Status	CPU status OK
CPU Temp	CPU temperature
Power Status	Power on/off
Batt Amp	Battery amp
Batt Capacity	Battery capacity
Charger Amp	Battery charge amp
Charger Cycles	Battery charge cycle
Charger Volt	Battery charge voltage
Batt Temp	Battery temperature
Batt Run Time	Battery run time Note: The duration of data preservation indicated by the Batt Run Time sensor is an estimate. Do not rely on its exact value.
Batt 8.0V	Battery 8.0 voltage
NVMEM 1.8V	NVMEM 1.8 voltage
NVMEM 8.0V	NVMEM 8.0 voltage
SAS Status	SAS status OK

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



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

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

You use the Data ONTAP CLI to administer the storage system and the BMC.

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

Next topics

[What the BMC does](#) on page 247

[Ways to configure the BMC](#) on page 248

[How to manage the BMC with Data ONTAP](#) on page 251

[How to log in to the BMC](#) on page 253

[How to manage the storage system with the BMC](#) on page 255

[How to display information about the storage system and the BMC](#) on page 261

[Comparison of Data ONTAP and BMC commands](#) on page 266

[How to troubleshoot the storage system with the BMC](#) on page 268

[How to update the BMC firmware](#) on page 268

[How to troubleshoot BMC problems](#) on page 269

What the BMC does

The BMC command line interface (CLI) commands enable you to remotely access and administer the storage system and diagnose error conditions. Also, the BMC extends AutoSupport capabilities by sending alerts and notifications through an AutoSupport message.

The BMC provides the following remote management capabilities for the storage system. You use the BMC CLI commands to perform the following tasks:

- Administer the storage system using the Data ONTAP CLI by using the BMC's system console redirection feature
- 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 BMC's system console log
 - View storage system events, captured in the BMC's system event log
 - Initiate a storage system core dump
 - Power-cycle the storage system (or turn it on or off)
- Monitor environmental and non-environmental sensors for the controller module and the NVMEM battery.
- Switch between the primary and the backup firmware hubs to assist in bootup and recovery from a corrupted image in the storage system's primary firmware hub.

The BMC extends AutoSupport capabilities by sending alerts and “down system” or “down filer” notifications through an AutoSupport message when the storage system goes down, regardless of whether the storage system can send AutoSupport messages. Other than generating these messages on behalf of a system that is down, and attaching additional diagnostic information to AutoSupport messages, the BMC has no effect on the storage system's AutoSupport functionality. The system's

AutoSupport behavior is the same as it would be without BMC installed. The AutoSupport configuration settings and message content behavior of the BMC are inherited from Data ONTAP.

The BMC supports the SSH protocol for CLI access from UNIX clients and PuTTY for CLI access from PC clients. Telnet and RSH are not supported. These protocols are not available on the BMC, and system options to enable or disable them have no effect on the BMC.

Note: The BMC ignores the `autologout.telnet.timeout` and `autologout.console.timeout` options. The settings for these options do not have any effect on the BMC.

Related concepts

[How to troubleshoot the storage system with the BMC](#) on page 268

[The AutoSupport feature](#) on page 175

Ways to configure the BMC

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

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

- Initializing a storage system that has the BMC
When the storage system setup process is complete, the `bmc setup` command runs automatically. For more information about the entire setup process, see the *Data ONTAP 7-Mode Software Setup Guide*.
- Running the Data ONTAP `setup` script
The `setup` script ends by initiating the `bmc setup` command.
- Running the Data ONTAP `bmc setup` command

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

Next topics

[Prerequisites for configuring the BMC](#) on page 248

[Configuring the BMC](#) on page 249

Prerequisites for configuring the BMC

Before you configure the BMC, you need to gather information about your network and your AutoSupport settings.

The following is the information you need to gather:

- Network information
You can configure the BMC using DHCP or static addressing.

- If you are using DHCP addressing, you need the BMC's MAC address. You can obtain it by using the `bmc status` command or from the MAC address label on the BMC.

Note: If you do not provide a valid BMC MAC address, an EMS message shows up to remind you during system bootup or when you use the `bmc status` or the `setup` command.

- If you are using a static IP address, you need the following information:
 - An available static IP address
 - The netmask of your network
 - The gateway of your network
- AutoSupport settings

The BMC uses the same mailhost information that Data ONTAP uses for AutoSupport. The BMC does not have its own mailhost setting. The BMC sends event notifications based on the following Data ONTAP AutoSupport settings:

 - `autosupport.to`
 - `autosupport.mailhost`

It is best that you set at least the `autosupport.to` option before configuring the BMC.

Note: The BMC does not rely on the storage system's `autosupport.support.transport` option to send notifications. The BMC uses the Simple Mail Transport Protocol (SMTP).

Related tasks

[Configuring AutoSupport](#) on page 177

Configuring the BMC

You can use the `setup` command or the `bmc setup` command to configure the BMC.

It is best to configure AutoSupport before configuring the BMC. Data ONTAP automatically sends AutoSupport configuration to the BMC, allowing the BMC to send alerts and notifications through an AutoSupport message.

Steps

1. At the storage system prompt, enter one of the following commands:
 - `setup`
 - `bmc setup`

If you enter `setup`, the `bmc setup` script starts automatically after the `setup` command runs.

2. When the BMC setup asks you whether to configure the BMC, enter `y`.
3. Do one of the following when the BMC setup asks you whether to enable DHCP on the BMC.

- To use DHCP addressing, enter **y**.
- To use static addressing, enter **n**.

Note: DHCPv6 servers are not currently supported.

4. If you do not enable DHCP for the BMC, the BMC setup prompts you for static IP information. Provide the following information when prompted:

- The IP address for the BMC
- The netmask for the BMC
- The IP address for the BMC gateway
- The name or IP address of the mail host to use for AutoSupport

Note: Currently, you can use only IPv4 addresses to connect to the BMC.

5. Enter the Address Resolution Protocol (ARP) interval for the BMC when you are prompted.
6. If the BMC setup prompts you to reboot the system, enter the following command at the storage system prompt:

```
reboot
```

7. At the storage system prompt, enter the following command to verify that the BMC's network configuration is correct:

```
bmc status
```

8. At the storage system prompt, enter the following command to verify that the BMC AutoSupport function is working properly:

```
bmc test autosupport
```

Note: The BMC uses the same mail host information that Data ONTAP uses for AutoSupport. The `bmc test autosupport` command requires that you set up the `autosupport.to` option properly.

You have successfully set up the BMC AutoSupport function when the system displays the following output:

Please check ASUP message on your recipient mailbox.

Examples for configuring the BMC and displaying configuration information

The following example shows how the BMC is configured to use a static IP address:

```
The Baseboard Management Controller (BMC) provides remote management capabilities including console redirection, logging and power control. It also extends autosupport by sending down filer event alerts.
```

```
Would you like to configure the BMC [y]: y
Would you like to enable DHCP on the BMC LAN interface [y]: n
Please enter the IP address for the BMC []: 10.98.148.61
Please enter the netmask for the BMC []: 255.255.255.0
Please enter the IP address for the BMC Gateway []: 10.98.148.1
Please enter gratuitous ARP Interval for the BMC [10 sec (max 60)]:
```

```
The mail host is required by your system to enable BMC to send
ASUP message when filer is down
```

```
Please enter the name or IP address of the mail host []:
You may use the autosupport options to configure alert destinations.
Now type 'reboot' for changes to take effect.
```

The following example shows how the BMC is configured to use DHCP:

```
The Baseboard Management Controller (BMC) provides remote management
capabilities including console redirection, logging and power control.
It also extends autosupport by sending down filer event alerts.
```

```
Would you like to configure the BMC [y]: y
Would you like to enable DHCP on the BMC LAN interface [y]: y
Please enter gratuitous ARP Interval for the BMC [10 sec (max 60)]:
```

```
The mail host is required by your system to enable BMC to send
ASUP message when filer is down
```

```
Please enter the name or IP address of the mail host:
You may use the autosupport options to configure alert destinations.
Now type 'reboot' for changes to take effect.
```

The following example displays the BMC status and configuration information:

```
Baseboard Management Controller:
Firmware Version: 1.0
IPMI version: 2.0
DHCP: off
BMC MAC address: ff:ff:ff:ff:ff:ff
IP address: 10.98.148.61
IP mask: 255.255.255.0
Gateway IP address: 10.98.148.1
BMC ARP interval: 10 seconds
BMC has (1) user: naroot
ASUP enabled: on
ASUP mailhost: mailhost@companyname.com
ASUP from: postmaster@companyname.com
ASUP recipients: recipient@companyname.com
Uptime: 0 Days, 04:47:45
```

Related concepts

[Prerequisites for configuring the BMC](#) on page 248

How to manage the BMC with Data ONTAP

You can manage the BMC from the storage system by using the Data ONTAP `bmc` commands and by changing the AutoSupport settings that are used by the BMC.

Next topics

[Data ONTAP commands for the BMC](#) on page 252

[BMC and AutoSupport options](#) on page 253

Data ONTAP commands for the BMC

Data ONTAP provides `bmc` commands that allow you to manage the BMC, including setting up the BMC, rebooting the BMC, displaying the status of the BMC, and updating the BMC firmware.

The following table describes the Data ONTAP commands for the BMC. These commands are also described in the `na_bmc(1)` 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 BMC. This is normal behavior.

Data ONTAP Command for the BMC	
<code>bmc help</code>	Displays the list of <code>bmc</code> commands available with the current release of Data ONTAP.
<code>bmc setup</code>	Initiates the interactive BMC setup program to configure the LAN settings.
<code>bmc status</code>	Displays BMC status. Note: The Data ONTAP <code>sysconfig</code> command displays both the storage system and the BMC status.
<code>bmc test autosupport</code>	Sends a test e-mail to all recipients specified with these options: <ul style="list-style-type: none"> • <code>autosupport.enable</code> • <code>autosupport.support.enable</code> • <code>autosupport.mailhost</code> • <code>autosupport.from</code> • <code>autosupport.to</code> • <code>autosupport.noteto</code> • <code>autosupport.support.to</code>
<code>bmc reboot</code>	Reboots the BMC and causes the BMC to perform a self-test. Any console connection through the BMC is lost. Note: Upon a BMC reboot, the console connection through the BMC is briefly interrupted. The console window may freeze for a few seconds.

Related concepts

[Ways to configure the BMC](#) on page 248

BMC and AutoSupport options

The BMC extends AutoSupport capabilities by sending alerts and notifications through an AutoSupport message. You can manage AutoSupport event messages and alerts by using the autosupport options.

The BMC uses the settings of the following options to send event messages and alerts:

- `autosupport.to`
- `autosupport.mailhost`

You use the `autosupport.content` option to change the amount of information displayed by Data ONTAP and BMC AutoSupport commands:

Related concepts

[AutoSupport options](#) on page 177

[Contents of AutoSupport event messages](#) on page 186

How to log in to the BMC

To log in to the BMC, you must install a Secure Shell client application and ensure that you have administrative privileges on the storage system.

The following are the prerequisites for logging in to the BMC:

- A Secure Shell client application that is appropriate for your administration host, such as SSH, or OpenSSH for UNIX hosts, or PuTTY for Windows hosts
- The root, naroot, or Administrator account

The password for all three account names is the same as the Data ONTAP root password.

Note: The BMC uses the Data ONTAP root password to allow access over the LAN with SSH. To access the BMC by using SSH, you must configure the Data ONTAP root password. BMC accepts passwords that are no more than 16 characters.

Next topics

[Accessing the BMC from a console](#) on page 254

[Logging in to the BMC from a UNIX host](#) on page 254

[Logging in to the BMC from a Windows host](#) on page 254

[BMC CLI and system console sessions](#) on page 255

Related concepts

[How to manage administrator and diagnostic access](#) on page 115

Accessing the BMC from a console

You can access the BMC from a console that is attached by a cable to the system's serial port.

Step

1. Press Ctrl-G at the storage system prompt.

Note: To return to the console from the BMC, enter `system console` at the BMC prompt.

The BMC prompt appears.

Logging in to the BMC from a UNIX host

You can log in to the BMC from a UNIX host, if a secure shell application is installed on the UNIX host.

Ensure that a secure shell application is installed on the UNIX host.

Step

1. Enter the following command at the UNIX host prompt:

```
secure_shell_app username@BMC_IP_address
```

username can be root, naroot, or Administrator.

Example

The following example shows how to log in to the BMC as naroot:

```
ssh naroot@192.0.2.123
```

Logging in to the BMC from a Windows host

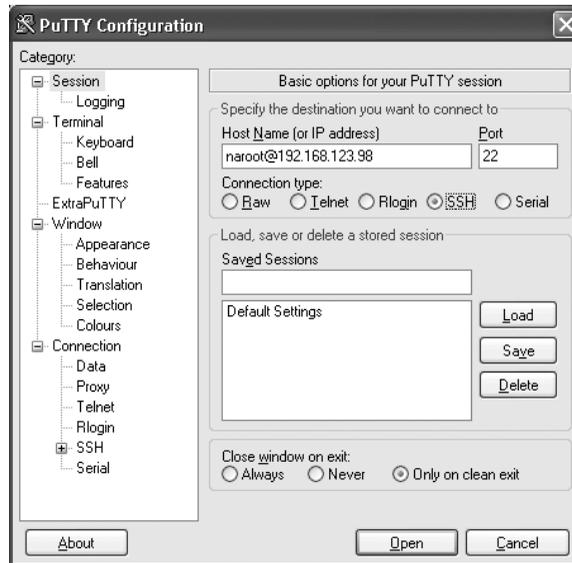
You can log in to the BMC from a Windows host if a Secure Shell application for Windows, such as PuTTY, is installed.

Steps

1. Start a Windows session for the Secure Shell application.

Example

You can use the PuTTY Configuration window to log in to the BMC:



2. Enter the IP address of the BMC.
3. Ensure that the SSH protocol is selected.
4. Click **Open**.
5. When you are prompted, use the root, naroot, or Administrator account to log in to the BMC.

The password for all three user accounts is the same as the Data ONTAP root password.

BMC CLI and system console sessions

Only one administrator can be logged in to an active BMC CLI session at a time. However, the BMC allows you to open both a BMC CLI session and a separate, BMC-redirected system console session simultaneously.

When you use the BMC CLI to start a system console session, the BMC CLI is suspended, and the system console session is started. When you exit the system console session, the BMC CLI session resumes.

The BMC prompt is displayed as `bmc shell ->`. For example, if your storage system is named `toaster`, the storage system prompt is `toaster>` and the prompt for the BMC session is `bmc shell ->`.

How to manage the storage system with the BMC

The BMC enables you to manage the storage system by using the BMC CLI. The BMC CLI has the same features available in the Data ONTAP CLI.

The CLI features include:

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

Like the Data ONTAP CLI, the BMC CLI provides two privilege levels, admin and advanced, with different command sets.

Note: The BMC CLI commands are *not* documented in online command-line manual (man) pages.

Next topics

[Online help at the BMC CLI](#) on page 256

[What you can do in BMC admin mode](#) on page 257

[BMC admin mode command syntax summary](#) on page 257

[What you can do in BMC advanced mode](#) on page 259

[Connecting to the storage system console from the BMC](#) on page 260

[Managing the controller module power of the storage system from the BMC](#) on page 260

Related concepts

[Data ONTAP command-line interface](#) on page 37

[Data ONTAP commands at different privilege levels](#) on page 39

Online help at the BMC CLI

The BMC Help displays all the available BMC commands when you enter the question mark (?) or help at the BMC prompt.

The following example shows the BMC CLI Help:

```
bmc shell -> ?
exit
bmc config
bmc config autoneg [enabled|disabled]
bmc config dhcp [on|off]
bmc config duplex [full|half]
bmc config gateway [gateway]
...
```

If a command has subcommands, you can see them by entering the command name after the help command, as shown in the following example:

```
bmc shell -> help events
events all           Print all system events
events info         Print SEL(system event
log)                information
events latest [N]   Print N latest system events
events oldest [N]   Print N oldest system events
events search [attr=N] Search for events
```

by	attribute/value pair
events show [N]	Print event N

What you can do in BMC admin mode

In the BMC admin mode, you can use the BMC commands to perform most tasks.

In admin mode, you can use the BMC commands to perform the following tasks:

- Connect to the storage system console (`system console`)
- Control the storage system power (`system power {on | off | cycle}`)
- Display the following information:
 - Available commands (`help` or `?`)
 - Syntax usage for a specific command (`help command`)
 - Storage system information (`system show`)
 - Storage system power status (`system power status`)
 - Storage system console logs (`system log`)
 - System hardware sensors and their status (`sensors subcommand`)
 - Chassis FRU information (`fru show`)
 - Events that occur on the storage system (`events subcommand`)
 - Current privilege level (`priv`)
 - BMC configuration information (`bmc config`)
 - BMC version (`bmc show`)
- Dump the storage system core and reset the storage system (`system core`)
- Exit from the BMC CLI (`exit`)
- Configure BMC (`bmc config subcommand`)
- Set the user mode privilege level (`priv set [admin | advanced]`)

BMC admin mode command syntax summary

The BMC commands in admin mode enable you to perform most of the tasks supported by the BMC.

The following table provides a quick reference of the command syntax for the BMC commands you can use in admin mode.

BMC admin mode command syntax	Description
<code>help [command]</code>	Displays a list of available commands. If a command is specified, displays the subcommands available for that command or its syntax usage.
<code>exit</code>	Exits from the BMC command line interface.
<code>bmc config</code>	Displays the BMC configuration.

BMC admin mode command syntax	Description
<pre>bmc config {autoneg [enabled disabled] dhcp [on off] duplex [full half] gateway [<i>gateway</i>] ipaddr [<i>ip-address</i>] netmask [<i>netmask</i>] speed [10 100]}</pre>	<p>Enables or disables Ethernet port auto negotiation. Sets BMC DHCP, Ethernet port duplex mode, BMC IP gateway, BMC IP address, BMC IP netmask, or Ethernet port speed at 10M or 100M.</p>
<pre>bmc show</pre>	<p>Displays BMC version and system information.</p>
<pre>events {all info latest [<i>M</i>] oldest [<i>N</i>] search [attr=<i>N</i>] show [<i>N</i>]}</pre>	<p>Displays storage system events logged by the BMC, including all system events, system event log (SEL) information, <i>N</i> latest system events, <i>N</i> oldest system events, events by attribute/value pair, or event <i>N</i>.</p> <p>For example, the following command displays events of the sensor whose ID is #d1:</p> <pre>events search id=#d1</pre> <p>Note: You can find the sensor ID by using <code>sensors show</code>. Use <code>id=#ff</code> for Data ONTAP and BMC status events.</p>
<pre>fru show</pre>	<p>Displays chassis FRU information.</p>
<pre>priv</pre>	<p>Displays current privilege level.</p>
<pre>priv set [admin advanced]</pre>	<p>Sets the privilege level to access the specified mode. The default is the <code>admin</code> mode.</p>
<pre>sensors show</pre>	<p>Displays current state of sensors.</p>
<pre>sensors search [attr=<i>N</i>]</pre>	<p>Searches a sensor by its ID.</p> <p>For example, the following command displays current state of sensor #09.</p> <pre>sensors search id=#09</pre> <p>Note: You can find the sensor ID by using <code>sensors show</code>.</p>
<pre>system console</pre>	<p>Logs in to Data ONTAP CLI.</p> <p>Note: Use Ctrl-G to return to the BMC prompt.</p>
<pre>system core</pre>	<p>Dumps the storage system core and resets the storage system. This command has the same effect as pressing the Non-maskable Interrupt (NMI) button on a storage system.</p> <p>Note: The BMC stays operational as long as input power to the storage system is not interrupted.</p>
<pre>system log</pre>	<p>Displays the system console history.</p>

BMC admin mode command syntax	Description
<code>system power {on off cycle status}</code>	<p>Turns the storage system on or off, performs a power cycle, or displays the power status. Standby power stays on. Using the <code>system power</code> command may cause an improper shutdown of the storage system. During power-cycling, there is a brief pause before power is turned back on.</p> <p>Note: If a storage system is power-cycled with the system in a high-availability configuration, the other system takes over and the rebooted system comes up in a “waiting for giveback” mode.</p>
<code>system show</code>	<p>Displays system information.</p> <p>Note: The <code>fault</code> field of the output displays system status, which may be <code>none</code> (no fault), <code>pcm</code> (the Processor Controller Module has a fault), or <code>system</code> (Data ONTAP has detected a system level fault that does not involve the PCM).</p>

What you can do in BMC advanced mode

The BMC advanced commands display more information than is available in administrative mode, including active BMC network services, storage system events logged by the BMC, and BMC battery information.

The following table provides a quick reference of the command syntax for the additional BMC commands that you can use in advanced mode.

BMC advanced mode command syntax	Description
<code>battery {show flash}</code>	<p>Displays BMC battery information or initiating a forced update of the battery firmware.</p> <p>You use <code>battery flash</code> if the automatic battery firmware upgrade process has failed for some reason.</p>
<code>events clear</code>	Deletes all storage system events logged by the BMC.
<code>service info</code>	Displays active BMC network services.
<code>system reset [primary backup]</code>	<p>Reboots the storage system using primary or backup firmware.</p> <p>Note: You use the <code>system show</code> command in BMC mode to display the firmware in use. You use the backup firmware to reboot the storage system if the primary firmware is corrupted and cannot be used for booting.</p>

Attention: Advanced commands are potentially dangerous. Use them only when directed to do so by your technical support personnel.

Connecting to the storage system console from the BMC

The BMC `system console` command enables you to log in to the storage system from the BMC.

Steps

1. Enter the following command at the BMC prompt:

```
system console
```

The message “Press ^G to enter BMC command shell” appears.

2. Press Enter to see the storage system prompt.

You use Ctrl-G to exit from the storage system console and return to the BMC CLI.

The storage system prompt appears, and you can enter Data ONTAP commands.

Example of connecting to the storage system console from the BMC

The following example shows the result of entering the `system console` command at the BMC prompt. The `vol status` command is entered at the storage system prompt, followed by Ctrl-G, which returns you to the BMC prompt.

```
bmc shell -> system console
Press ^G to enter BMC command shell
```

(Press Enter to see the storage system prompt.)

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

(Information about all of the existing volumes is displayed.)

(Press Ctrl-G to exit from the storage system console and return to the BMC prompt.)

```
bmc shell ->
```

Related tasks

[Using the remote management device to access the system console](#) on page 48

Managing the controller module power of the storage system from the BMC

The BMC's `system power` command enables you to turn the power on or off or to power-cycle the storage system remotely.

The `system power cycle` command automatically turns system power off and then back on. The power supplies provide a standby voltage that is always present, even when the storage system is off.

This keeps the BMC running without interruption. However, turning the power off or power-cycling the storage system may cause an improper shutdown of the storage system (also called a dirty shutdown).

Steps

1. Enter the following command at the BMC prompt:

```
system power cycle
```

Example

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

A warning message indicates that issuing the `system power` command is not a substitute for a graceful shutdown using the Data ONTAP `halt` command.

2. To turn off the storage system, enter `y`.

A few seconds later, the storage system is turned back on, and the boot environment prompt appears. In the meantime, the BMC prompt awaits your next command.

How to display information about the storage system and the BMC

The BMC provides several ways to display information about the storage system and the BMC itself. You can display real-time information using the commands in admin or advanced mode, or you can display information stored in the BMC's system event log (SEL) or console log.

You can also view the information displayed in the AutoSupport messages generated by the BMC. Most of the information is stored in the SEL or in captured console messages.

All log entries are recorded with Coordinated Universal Time (UTC) for the time format.

Note: The BMC does not use the time zone setting from the storage system.

Next topics

[BMC CLI commands that display real-time information](#) on page 261

[System event log and the BMC](#) on page 264

[System console log and the BMC](#) on page 265

[AutoSupport messages for systems with the BMC](#) on page 265

BMC CLI commands that display real-time information

Using the BMC CLI commands in admin mode, you can view information such as the BMC version and configuration, system console log history, storage system events, and storage system power status. Using the BMC CLI commands in advanced mode, you can view information about the BMC battery and active BMC network services.

Using the BMC CLI commands in admin mode, you can view the following information:

- The storage system information (`system show`)
- BMC version and system information (`bmc show`)
- BMC configuration information (`bmc config`)
- The state of BMC sensors (`sensors show`)

Note: In the `sensors show` output, the duration of data preservation indicated by the Batt Run Time sensor is an estimate. Do not rely on its exact value.

- The system console log history (`system log`)
- Storage system events logged by the BMC (`events {all | info | latest [N] | oldest [N] | search [attr=N] | show [N]}`)
- The storage system power status (`system power status`)
- Chassis FRU information (`fru show`)
- Current privilege level (`priv`)

Using the BMC CLI commands in advanced mode, you can view the following information:

- BMC battery information (`battery show`)
- Active BMC network services (`service info`)

BMC CLI in admin mode

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

```
bmc shell -> system show
power          :on
fault          :none
reset          :off
name           :SystemName
product        :ModelNumber
serial-number  :1070065
firmware       :primary
mellanox       :normal
nvmem          :enabled
```

```
bmc shell -> bmc show
FirmwareVersion:1.0X8
SystemUptime   :7 Days, 10:15:20
Date           :03/29/2007 21:35:10 GMT
```

```
bmc shell -> bmc config
ipaddr         :10.98.148.61
netmask        :255.255.255.0
gateway        :10.98.148.1
mac            :00:a0:98:05:2d:1e
dhcp           :off
link           :up
autoneg        :enabled
```

```
speed      :100
duplex     :full
```

```
bmc shell -> sensors show
```

Name	State	Id	Reading	Crit-Low	Warn-Lo	Warn-Hi	Crit-Hi
1.1V	Normal	#77	1121 mV	955 mV	--	--	1239 mV
1.2V	Normal	#76	1239 mV	1038 mV	--	--	1357 mV
1.5V	Normal	#75	1522 mV	1309 mV	--	--	1699 mV
1.8V	Normal	#74	1829 mV	1569 mV	--	--	2029 mV
12.0V	Normal	#70	12080 mV	10160 mV	--	--	13840 mV
2.5V	Normal	#73	2539 mV	2116 mV	--	--	2870 mV
3.3V	Normal	#72	3374 mV	2808 mV	--	--	3799 mV
BIOS Status	Normal	#f0	System #2f	--	--	--	--
Batt 8.0V	Normal	#50	7872 mV	--	--	8512 mV	8576 mV
Batt Amp	Normal	#59	0 mA	--	--	2112 mA	2208 mA
Batt Capacity	Normal	#54	3744 mAh	--	--	--	--
Batt Run Time	Normal	#55	182 h	72 h	80 h	--	--
Batt Temp	Normal	#51	33 C	0 C	10 C	45 C	60 C
Board Temp Bot	Normal	#08	52 C	-3 C	7 C	69 C	79 C
Board Temp Top	Normal	#07	40 C	-3 C	7 C	54 C	62 C
CPU 1.2V	Normal	#71	1180 mV	1038 mV	--	--	1357 mV
CPU Status	Normal	#f1	Ok	--	--	--	--
CPU Temp	Normal	#09	63 C	--	--	--	126 C
Charger Amp	Normal	#53	0 mA	--	--	--	--
Charger Cycles	Normal	#58	4	--	--	250	251
Charger Volt	Normal	#52	8192 mV	--	--	--	--
NVMEM 1.8V	Normal	#0b	1790 mV	1621 mV	1706 mV	1889 mV	1974 mV
NVMEM 8.0V	Normal	#0a	7648 mV	--	--	8508 mV	8604 mV
Power Status	Normal	#d1	Power On	--	--	--	--
SAS Status	Normal	#b6	Ok	--	--	--	--

```
bmc shell -> system power status
```

```
power      :on
```

```
bmc shell -> fru show
```

```
board_mfg      :CompanyName
board_product  :111-00238+P2A
board_serial   :1070065
board_part     :110-00038+P2A
product_mfg    :CompanyName
product_name   :ProductModel
product_part   :
product_version:
product_serial :1070065
system_serial  :0041070065
```

```
bmc shell -> priv
```

```
admin
```

The BMC CLI in advanced mode

The following examples show how information is displayed using commands at the BMC advanced mode prompt:

```
bmc shell*-> battery show
chemistry      :LION
device-name    :bq20z80
expected-load-mw:162
id             :27100011
manufacturer   :AVT
manufacture-date:6/28/2006
rev_cell       :2
rev_firmware   :200
rev_hardware    :c0
serial         :80b6
status         :full
test-capacity  :disabled
```

```
bmc shell*-> service info
ssh           :enabled
rmcp         :disabled
```

System event log and the BMC

The BMC has a nonvolatile memory buffer that stores up to 512 system events in a system event log (SEL). The SEL is stored in onboard flash memory on the BMC.

The SEL stores each audit log entry as an audit event. You can view these audit log entries, along with other stored events, by using the BMC `events` commands. You can also use the `events search` command to perform a quick search for audit logs from the SEL.

The SEL stores platform-specific events. This log is self-contained and does not support the Syslog Translator.

The primary purpose of the SEL is to help you diagnose system issues. The event list from the SEL is automatically sent by the BMC to specified recipients through an AutoSupport message.

The records contain the following data:

- Hardware events detected by the BMC—for example, system sensor status about power supplies, voltage, or other components
- Errors (generated by the storage system or the BMC) detected by the BMC—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 BMC by the storage system—for example, a system panic, a communication failure, an unexpected boot environment prompt, a 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 BMC’s clock to time-stamp events. BMC 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 'pre-init time'.

The following example shows the result of entering BMC `events` command:

```
bmc shell -> events search id=#dl
Event TimeStamp          Id  Sensor          Description
```

```
-----
42    03/30/2007 16:29:53 GMT #d1 Power Status    Power Off
43    03/30/2007 16:30:04 GMT #d1 Power Status    Power On
Total Entries=2
```

System console log and the BMC

The BMC monitors the storage system console regardless of whether administrators are logged in or connected to the console. When storage system messages are sent to the console, the BMC stores them in the system console log, which resides in a 64-KB buffer in its main memory.

The system console log can store approximately 1,000 lines of system console messages. When the buffer is full, the oldest messages are overwritten by the newest messages.

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

When a “down filer,” a system hang, or a reboot loop condition occurs, the system console logs are attached to the AutoSupport messages sent by the BMC, regardless of the state of the `autosupport.content` option.

You display the contents of the system console log with the BMC CLI `system log` command.

Note: Entering the BMC CLI command `system log` is only recommended from the SSH interface, because the 9600 baud serial console interface is very slow to display the entire log.

AutoSupport messages for systems with the BMC

For storage systems with the BMC, there are two additional types of AutoSupport messages—BMC-generated AutoSupport messages about the storage system, and storage system-generated AutoSupport messages about the BMC.

BMC-generated AutoSupport messages include the following information:

- In the subject line—A system notification from the BMC of the storage system, listing the system condition or event that caused the AutoSupport message, and the log level.
- In the message body—The BMC configuration and version information, the storage system ID, serial number, model and host name.
- In the attachments—the system event logs, the system sensor state as determined by the BMC, and system console logs.

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

- The storage system reboots unexpectedly.
- The storage system stops communicating with the BMC.
- A watchdog reset occurs.

The watchdog is a built-in hardware sensor that monitors the storage system for a hung or unresponsive condition. If the watchdog detects this condition, it resets the storage system so that

the system can automatically reboot and begin functioning. This feature is sometimes called automatic server restart.

When the BMC detects a watchdog-generated event occurs on the storage system, it logs this event and, if needed, sends an AutoSupport alert for this event.

- The storage system is power-cycled.
- Firmware power-on self-test (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 BMC, a description of the BMC condition or event that caused the AutoSupport message, and the log level.
- 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 attachments—messages from `EMS`, `rc`, `exports`, `hosts`, `resolv_conf`, `nsswitch_conf`, and `cm_stats`

Typical storage system AutoSupport messages about the BMC occur under the following conditions:

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

Related concepts

[AutoSupport messages](#) on page 185

Comparison of Data ONTAP and BMC commands

Whether you use a Data ONTAP command or a BMC command to manage the BMC depends on the task you want to perform.

The following table shows the comparison of Data ONTAP and BMC commands.

Action	Data ONTAP Command or Procedure	BMC Command
Set up BMC in a new storage system	<p>Turn on the new storage system.</p> <p>During bootup, press Ctrl-C when prompted to access the boot menu.</p> <p>From the menu, select 4 to initialize disks and set up a root volume.</p> <p>Note: After the Data ONTAP setup script is completed, the BMC <code>bmc setup</code> script is initiated.</p>	

Action	Data ONTAP Command or Procedure	BMC Command
Reconfigure a BMC in an existing storage system	setup Note: After the Data ONTAP setup script is completed, the bmc setup script is initiated.	
Initiate the bmc setup script to configure BMC	bmc setup	
Test the BMC's AutoSupport setting	bmc test autosupport	
Display Data ONTAP bmc commands	bmc help	
Log in to the BMC	From the system console, press Ctrl-G.	From a UNIX host, enter the following command: ssh naroot@BMC_IP_addr
Display BMC CLI commands		help or ?
Display the twenty most recent events logged by BMC		events latest 20
Display the five oldest events logged by BMC		events oldest 5
Display a summary of information about the records in the events log		events info
Display the BMC configuration and version information	bmc status or sysconfig -v Note: sysconfig -v requires advanced mode.	bmc config or bmc show
Display the list and the status of system hardware sensors		sensors show or sensors search [attr=N]
Log in to the system to manage storage system resources		system console Note: Use exit to exit the BMC shell.
Display what has occurred at the storage system console		system log
Dump the system core and reset the storage system		system core
Control the storage system power or display the power status		system power {on off cycle status}

Action	Data ONTAP Command or Procedure	BMC Command
Reboot the storage system	reboot	system reset [primary backup] Note: Requires advanced mode.
Reset the BMC	bmc reboot	

Attention: Advanced commands are potentially dangerous. Use them only when directed to do so by your technical support personnel.

How to troubleshoot the storage system with the BMC

When you encounter a problem with the storage system, you can use the BMC 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.

If this condition occurs...	And you want to...	Enter this command at the BMC 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 show [N] events latest [N]
	Search for specific events in the SEL	events search [attr=N]
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 Note: Requires advanced mode.

How to update the BMC firmware

The BMC firmware is bundled with the Data ONTAP distribution, which is stored on the boot device, such as a PC CompactFlash card. The `update_bmc` command, entered at the boot

environment prompt, updates the BMC firmware from the Data ONTAP image on the boot device. The BMC firmware is also available for download.

For instructions on how to download and update the BMC firmware, see the *Data ONTAP 7-Mode Upgrade Guide*.

How to troubleshoot BMC problems

BMC problems might result from communication problems, configuration problems, connection problems, BMC hardware failures, or BMC firmware update problems.

Next topics

[Troubleshooting BMC communication problems](#) on page 269

[Troubleshooting BMC configuration problems](#) on page 269

[Troubleshooting BMC connection problems](#) on page 270

[Troubleshooting BMC hardware failures](#) on page 270

[Troubleshooting BMC firmware update problems](#) on page 272

Troubleshooting BMC communication problems

A communication failure between the storage system and the BMC might result in BMC problems.

Step

1. If there is a communication failure between the storage system and the BMC, search for EMS events titled:
`[bmc.orftp.failed:warning]: BMC communication error, (reason)`

Troubleshooting BMC configuration problems

If you are having difficulty configuring the BMC, you should verify that the IP configuration is correct.

Steps

1. Verify the BMC is online and the IP configuration is correct by entering the following command at the storage system prompt:
`bmc status`
2. If the BMC is configured using DHCP, reconfigure the BMC using a static IP address by entering the following command at the storage system prompt:

```
bmc setup
```

Troubleshooting BMC connection problems

If you are having difficulty connecting to the BMC, you should verify that you are using a secure shell client and that the IP configuration is correct.

Steps

1. Verify that you are using a secure shell client to connect to the BMC.
2. From the storage system, verify the BMC is online and the IP configuration is correct by entering the following command at the storage system prompt:

```
bmc status
```

3. From the administration host, test the network connection for the BMC by entering the following command:

```
ping bmc_IP_address
```

4. If the ping fails, do one of the following:
 - Verify that the BMC 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 that the BMC has a valid IP address. At the storage system, enter the `bmc setup` command to use the DHCP server or assign a valid IP address.
 - Verify that the administration host has a route to the BMC.
5. From the storage system prompt, reboot the BMC by entering the following command:

```
bmc reboot
```

Note: It takes approximately one minute for the BMC to reboot.

6. If the BMC does not reboot, repeat Steps 2 through 5. If the BMC still does not reboot, contact technical support for assistance.

Related concepts

[How to log in to the BMC](#) on page 253

Troubleshooting BMC hardware failures

A BMC problem can occur when a hardware failure has occurred on the BMC.

When the BMC fails, an EMS event similar to the following can be found:

```
[asup.msg.bmc.heartbeat.stops:critical]: Data ONTAP lost communication with the baseboard management controller (BMC).
```

Steps

1. Run diagnostics by entering the following command from the boot environment prompt:

boot_diags

The diagnostics main menu appears.

- From the main menu, enter the following option:

mb

The motherboard diagnostic menu appears.

```

Enter Diag, Command or Option: mb
Motherboard Diagnostic
-----
1: Comprehensive motherboard diags  71: Show
PCI                                configuration
2: Misc. board test menu           72: Show detailed
PCI                                info
3: Cache test menu                 73: Initialize
real-                               time clock
4: On-board GbE test menu          75: System
serial                             info setup[Mfg]
5: On-board FCAL test menu
6: SAS Test Menu                   91: Enable/
disable                             looping
7: IB Test Menu                    92: Stop/
Continue                             looping on error
8: BMC Test Menu                   93: Extended/
Normal                               test mode
9: NVMEM Test Menu                 99: Exit

```

- From the diagnostic prompt, enter test number 8.

The BMC diagnostic menu appears.

```

Select test or feature by number [0]: 8
BMC Diagnostics
-----
1: Comprehensive Test              72: Get Reason
for                               Restart
2: BMC Self Test                  73: Show Device Info
3: Environment Test               74: Show SDR Info
4: SDR Read Test                  75: Show SEL Info
5: SEL Read Test                  76: Clear SEL [Mfg]
6: LCD Exercise                   77: Emergency
Shutdown                          [Mfg]
7: BMC Timer test                 78: BMC Update
Menu                               [Xtnd]
10: Show BMC SSH Keys             79: Dump SEL Records
                                  80: Dump Raw
                                  Records
SEL
41: BMC NMI Test
42: BMC Front Panel Button Test  91: Enable/
disable                             looping
43: SEL Write Test [Xtnd]        92: Stop/
continue                             on error
                                  93: Extended/

```

```
Normal                               test mode
71: Show BMC SEL Time                 99: Exit
```

4. Enter the appropriate test number from the diagnostic prompt. To perform a comprehensive test, enter test number 1.

Note: It takes several minutes to complete the comprehensive test.

The results of the test are displayed.

5. Based on the results of Step 4, diagnose the problem. If the problem persists, reset the BMC and repeat Steps 1 to 5.

If the problem still persists, contact technical support for assistance.

Troubleshooting BMC firmware update problems

A BMC firmware update failure can occur for a number of reasons. If a BMC firmware update fails, you may not be able to boot Data ONTAP.

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 BMC.
- The update failed when you attempted to install the new firmware at the BMC.
- The storage system was reset during the update.
- There was a power loss during update.

Steps

1. A/C power-cycle the storage system.
2. After the system is powered on again, update the BMC firmware by entering the following command from the boot environment prompt:

```
update_bmc
```

Note: If a previous power failure caused the system to boot from the backup firmware and halt at the boot environment prompt, repeat the `update_bmc` command at the boot environment prompt.

The `update_bmc` macro automatically updates the BMC firmware from the image on the boot device.

3. After the BMC firmware is updated, enter following command from the boot environment prompt to restart the system:

```
bye
```

4. Reboot the BMC by entering the following command at the storage system prompt:

```
bmc reboot
```

5. If the BMC still does not reboot, run diagnostics on the BMC.
6. If the BMC is not operational, contact technical support for assistance.

Note: If a BMC firmware update fails when Data ONTAP is running, Data ONTAP will try to recover the BMC by rebooting it. If the reboot fails, a message [asup.msg.bmc.heartbeat.stops:critical] is sent and the storage system is shut down. Contact technical support for assistance.

Related tasks

Troubleshooting BMC hardware failures on page 270

The Remote Support Agent as a firmware upgrade

The Remote Support Agent (RSA) is available as a firmware upgrade for storage systems that support the SP or the RLM.

Featuring remote data collection, intelligent core handling, and down controller notification, the RSA is an enhancement to the SP and the RLM firmware. The RSA enables NetApp to use the SP and the RLM for remote support. When problem diagnostics are needed, the RSA automatically uploads core files and transfers diagnostics data (such as log files) to NetApp technical support, reducing your involvement in the troubleshooting process.

For more information about the RSA, see the *Remote Support Agent Concepts and Customer Usage Guide* and the NetApp Remote Support Diagnostics Tool page on the NOW site.

Related information

NetApp Remote Support Diagnostics Tool page - <http://now.netapp.com/NOW/download/tools/rsa/>

System information

Data ONTAP enables you to display information about your storage system, including the system's configuration, storage components, aggregate and volume information, file statistics, environmental status, Fibre Channel information, and SAS adapter and expander information.

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

Next topics

[Getting storage system configuration information](#) on page 275

[Commands to display storage subsystem information](#) on page 277

[Getting aggregate information](#) on page 279

[Getting volume information](#) on page 280

[Getting a file statistics summary](#) on page 281

[Storage system environment information](#) on page 285

[Getting Fibre Channel information](#) on page 288

[Getting SAS adapter and expander information](#) on page 288

[Storage system information and the stats command](#) on page 289

[How to get system information using perfmon](#) on page 298

[How to get system information using perfstat](#) on page 298

Getting storage system configuration information

You can display configuration information about the storage system, including version information, hardware configuration, disk information, RAID and checksum information, tape drive information, volume information, and tape library information.

Step

1. Enter one of the following commands:

Command	Description
<code>version</code>	Displays the version of Data ONTAP currently running on a storage system.
<code>sysconfig</code>	Displays information about the storage system's hardware configuration. The exact types of information displayed depend on the command options.

Command	Description
sysconfig -c	<p>Checks that expansion cards are in the appropriate slots and reports any configuration errors.</p> <p>If there are no configuration errors, the <code>sysconfig -c</code> command reports the following: <code>sysconfig: There are no configuration errors.</code></p>
sysconfig -d	Displays product information about each disk in the storage system.
sysconfig -r	<p>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:</p> <ul style="list-style-type: none"> • 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 <p>Note: You can also obtain the information displayed by <code>sysconfig -r</code> from SNMP, using the custom Management Information Base (MIB). For information about SNMP, see the <i>Data ONTAP 7-Mode Network Management Guide</i>.</p>
sysconfig -t	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 <code>dump</code> and <code>restore</code> commands.
sysconfig -v	Displays RAID group and disk information about each traditional volume and aggregate.
sysconfig -m	Displays tape library information. Before you use this option, ensure that the storage system was booted with the <code>autoload</code> setting of the tape library off.
sysconfig -v	<p>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.</p> <p>Note: If you enter <code>sysconfig</code> without any options, information similar to what you get with <code>sysconfig -v</code> is displayed, but the information is abbreviated. When you report a problem to technical support, provide the information displayed by <code>sysconfig -v</code>. This information is useful for diagnosing system problems.</p>
sysconfig -a	Displays the same information as the <code>-v</code> option, but the information is more detailed.

Command	Description
<code>sysconfig -A</code>	<p>Displays storage system information gathered by the following commands, one after the other:</p> <ul style="list-style-type: none"> <code>sysconfig</code> <code>sysconfig -c</code> <code>sysconfig -d</code> <code>sysconfig -V</code> <code>sysconfig -r</code> <code>sysconfig -m</code> <p>Therefore, when you use the <code>sysconfig -A</code> command, Data ONTAP lists information about configuration errors, disk drives, medium changers, RAID details, tape devices, and aggregates.</p>

Note: You can also get system information, either interactively or with a script, using the `stats` command.

For more information about the `sysconfig` command, see the `na_sysconfig(1)` man page.

Related concepts

[Storage system information and the stats command](#) on page 289

Commands to display storage subsystem information

You can use the `acpadmin`, `environment`, `fcadmin`, `sasadmin`, `storage show`, and `sysconfig` commands to display information about your storage subsystems.

Note: For detailed information about these commands and their options, see the appropriate man pages.

Use this Data ONTAP command...	To display information about...
<code>acpadmin list_all</code>	Alternative Control Path (ACP) processors (DS4243 only).
<code>environment shelf</code>	Environmental information for each host adapter, including SES configuration and SES path.
<code>environment shelf_log</code>	Shelf-specific module log file information, for shelves that support this feature. Log information is sent to the <code>/etc/log/shelflog</code> directory and included as an attachment on AutoSupport reports.
<code>fcadmin channels</code>	WWPN information.
<code>fcadmin device_map</code>	What disks are on each loop and shelf.

Use this Data ONTAP command...	To display information about...
<code>fcadmin link_state</code>	How the ports are connected.
<code>sasadmin expander</code>	What disks are attached to expander PHYs.
<code>sasadmin expander_phy_state</code>	Expander PHY state, dongle state and event counters, PHY statistics.
<code>sasadmin shelf [short]</code>	The disks on each shelf (or a specific disk shelf), including a pictorial representation of disk placement (long or short view).
<code>storage show</code>	All disks and host adapters on the system.
<code>storage show acp</code>	Connectivity and status information for the Alternate Control Path (ACP) module (DS4243 only).
<code>storage show adapter</code>	FC host adapter attributes, including (as appropriate for the adapter type) a description, firmware revision level, Peripheral Component Interconnect (PCI) bus width, PCI clock speed, FC node name, cacheline size, FC packet size, link data rate, static random access memory (SRAM) parity, state, in use, redundant.
<code>storage show disk -p</code>	How many paths are available to each disk.
<code>storage show expander</code>	SAS expander attributes, including shelf name, channel, module, shelf ID, shelf UID, IOM state, and the following information for the disks attached to the expander: disk ID, port state, partial path timeout, link rate, invalid word count, running disparity count, PHY reset problem, CRC error count, and PHY change count.
<code>storage show hub</code>	Hub attributes: hub name, channel, loop, shelf ID, shelf user ID (UID), term switch, shelf state, ESH state, and hub activity for each disk ID: loop up count, invalid cyclic redundancy check (CRC) count, invalid word count, clock delta, insert count, stall count, util.
<code>storage show mc</code>	All media changer devices that are installed in the system.
<code>storage show port</code>	Switch ports connected to the system.
<code>storage show switch</code>	Switches connected to the system.
<code>storage show tape</code>	All tape drive devices attached to the system.

Use this Data ONTAP command...	To display information about...
<code>storage show tape supported [-v]</code>	All tape drives supported. With <code>-v</code> , information about density and compressions settings is also displayed.
<code>storage stats tape</code>	Statistics for all tape drives attached to the system.
<code>sysconfig -A</code>	All sysconfig reports, including configuration errors, disks, array LUNs, media changers, RAID details, tape devices, and aggregates.
<code>sysconfig -m</code>	Tape libraries.
<code>sysconfig -t</code>	Tape drives.

Getting aggregate information

You can display information about the configuration and the state of an aggregate.

About this task

You use the `aggr status` command to display information about aggregate configurations. The `aggr status` command works for aggregates that were created explicitly, as well as for the aggregates created automatically when traditional volumes were created. Because traditional volumes are tightly coupled with their containing aggregates, the `aggr status` command returns information for both aggregates and traditional volumes. In both cases, it is the aggregate information that is returned.

Step

1. Enter the following command:

```
aggr status [-d] [-r] [-v]
```

- With no options, the `aggr status` command displays a concise synopsis of aggregate states, including:
 - The aggregate name
 - Whether it is an aggregate (32-bit or a 64-bit) or traditional volume
 - Whether it is online, offline, or restricted
 - Its RAID type
 - Other states such as partial or degraded
 - Options that are enabled, either by default or through the `aggr options` or `vol options` command

Note: If you specify an aggregate, such as `aggr status aggr0`, the information for that aggregate is displayed. If you do not specify an aggregate, the status of all aggregates and traditional volumes in the storage system is displayed.

- The `-d` option displays information about disks.
The disk information is the same as the information from the `sysconfig -d` command.
- The `-r` option displays RAID, plex, and checksum information for an aggregate.
The display is the same as the `sysconfig -r` display.
- The `-v` option displays information about each RAID group within an aggregate or traditional volume, and the settings of the aggregate options.

Note: You can also get aggregate information, either interactively or with a script, using the `stats` command.

For more information about aggregates, see the *Data ONTAP 7-Mode Storage Management Guide*. For more information about the `aggr` command, see the `na_aggr(1)` man page.

Related concepts

[Storage system information and the stats command](#) on page 289

Getting volume information

You can display information about the configuration and the state of a volume.

Step

1. Enter the following command:

```
vol status [-d] [-r] [-v] [-l]
```

- With no options, the `vol status` command displays a concise synopsis of volume states, including:
 - Volume name
 - Whether it is a FlexVol or traditional volume
 - Whether it is online, offline, or restricted
 - Other status such as partial and degraded
 - Options that are enabled for the volume or its containing aggregate (through the `aggr options` or `vol options` command).

The `vol` command also displays RAID information for the volume's containing aggregate.

Note: If you specify a volume, such as `vol status vol0`, the information for that volume is displayed. If you do not specify a volume, the status of all volumes in the storage system is displayed.

- The `-d` option displays information about the volume's containing aggregate's disks.

The information displayed is the same as for the `sysconfig -d` command.

- The `-r` option displays RAID, plex, and checksum information for the volume's containing aggregate.

The information displayed is the same as for the `sysconfig -r` command.

- The `-v` option displays the state of all per-volume options and information about each plex and RAID group within the volume's containing aggregate.
- The `-l` option displays the language used by each volume.

Note: You can also get volume information, either interactively or with a script, using the `stats` command.

For more information about volumes, see the *Data ONTAP 7-Mode Storage Management Guide*. For more information about the `vol` command, see the `na_vol(1)` man page.

Related concepts

[Storage system information and the stats command](#) on page 289

Getting a file statistics summary

You can display 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 creation of Snapshot copies by enabling you to see when most file activity takes place on a volume. The information also helps you determine Snapshot copy disk consumption.

Step

1. Enter the following command:

```
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.
- `snapshot_name` is the name of the Snapshot copy.
- `volume_name` is the name of the volume.
- The `-g` option 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.
- The `-u` option 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.
- The `ages` option 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.

- 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 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

- The `sizes` option enables you to specify the breakdown of sizes, using a comma to separate each value. Default values are in bytes, but you can also use the following suffixes at the end of a number you specify:
 - K (kilobytes).
 - M (megabytes).
 - G (gigabytes).
 - * (a special value for listing all unique file sizes, one line per unique size). Using the * suffix can result in output of several thousands of lines.
- The `style` option controls the output style. The valid arguments are as follows:
 - `readable`—The default. This is what you see when you use the `filestats` command with no `style` option.
 - `table`—Use this argument when the `filestats` output will be used by processing programs.
 - `html`—Use this argument for output that will be read by a Web browser.
- 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.
- 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.

Result

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

Next topics

[Example of the filestats command with no options specified](#) on page 283

[Examples of the filestats command with ages option specified](#) on page 284

[Example of the filestats command with sizes option specified](#) on page 285

[Example of using the filestats command to determine volume capacity](#) on page 285

Example of the filestats command with no options specified

You can use the `filestats` command without any options to display information about a Snapshot copy, including a breakdown of files by size, age, user ID, and group ID, and the cumulative number of inodes for each value.

The following example shows sample output from the `filestats` command, without any options, for the `hourly.1` Snapshot copy on `vol0`.

```
toaster> 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.

Examples of the filestats command with ages option specified

You can use the `filestats` command with the `ages` option to display a daily breakdown of file changes in a volume.

The following example shows sample output from the `filestats` command with the `ages` option.

```
toaster> filestats ages 1D,2D,3D,4D,5D,6D,7D,8D,9D,10D,11D,12D,13D,14D
volume vol0 snapshot hourly.0
VOL=vol0 SNAPSHOT=hourly.0
INODES=1087338 COUNTED_INODES=7062 TOTAL_BYTES=3835561873 TOTAL_KB=3701388
```

FILE SIZE	CUMULATIVE COUNT	CUMULATIVE TOTAL KB
1K	2313	8428
10K	6057	30280
100K	6686	49148
1M	6949	167664
10M	7008	406648
100M	7053	1538644
1G	7062	3701388
MAX	7062	3701388

AGE (ATIME)	CUMULATIVE COUNT	CUMULATIVE TOTAL KB
1D	12	332
2D	20	364
3D	26	18016
4D	44	18208
5D	84	64984
6D	85	64984
7D	116	65308
8D	142	67552
9D	143	71620
10D	143	71620
11D	144	71624
12D	166	93216
13D	166	93216
14D	378	109712
MAX	7062	3701388

- You 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.
- You 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.

To display files with ages under 900 seconds (15 minutes), under 4 hours, and under 7 days, you use the following command:

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

The following example shows the age section of the output:

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

Example of the filestats command with sizes option specified

You can use the `filestats` command with the `sizes` option to specify the breakdown of sizes.

The following example shows the file size section of the output when `filestats sizes 500K, 2M, 1G volume vol0 snapshot hourly.1` is entered 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.

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

Example of using the filestats command to determine volume capacity

You can use the `filestats` command to determine when the most activity occurs on a volume during a given day so that you can effectively schedule creation of 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 creation of more Snapshot copies during these hours and fewer throughout the rest of the day. Scheduling creation of more Snapshot copies before or during increased file activity decreases the time between file changes and Snapshot copy creation.

For information about managing Snapshot copies, see the *Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide*.

Storage system environment information

You can display information about the storage system environment, including shelf status and temperature, storage system component information, storage system temperature, and devices attached to the storage system.

You use 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

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.

For more information about the `environment` command, see the `na_environment(1)` man page.

Next topics

[Getting environmental status information](#) on page 286

[Specifying a UPS device to be monitored](#) on page 287

[Enabling or disabling monitoring of UPS devices](#) on page 287

Getting environmental status information

The `environment` command enables you to display all environment information, shelf environment status, chassis environment status, and UPS devices information.

Step

1. Enter one of the following commands:

Command	Description
<code>environment status</code>	Displays all storage system environment information. Note: For systems that contain internal drives, the <code>environment status</code> command displays information for both the internal and the external storage environment.
<code>environment status shelf [adapter]</code>	Displays the shelf environmental status for all shelves if <code>adapter</code> is not specified. You use <code>adapter</code> to display shelf information for shelves attached to the specified adapter.

Command	Description
<code>environment chassis</code>	Displays the environmental status of all chassis components.
<code>environment chassis list-sensors</code>	Displays detailed information from all chassis sensors.
<code>ups status</code>	Displays the status of all UPS devices. You can add UPS devices to be monitored, enable or disable monitoring of UPS devices, or display the status of UPS devices.

Specifying a UPS device to be monitored

You can specify a UPS device to be monitored by the storage system's environmental monitoring software.

Step

1. Enter the following command:

```
ups add [-c community] IP_address
```

- You use `-c community` 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

You can enable or disable monitoring of one or more UPS devices.

Step

1. Enter the following command:

```
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.

Getting Fibre Channel information

You can display Fibre Channel information such as the link statistics for all disks on a loop, internal Fibre Channel driver statistics, and the relative physical positions of drives on a loop.

Step

1. To display Fibre Channel information, enter one of the following commands:

Command	Description
<code>fcstat</code> <code>link_stats</code>	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.
<code>fcstat</code> <code>fcsl_stats</code>	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.
<code>fcstat</code> <code>device_map</code>	Displays the relative physical positions of drives on a loop and the mapping of devices to disk shelves.

Note: You can also get Fiber Channel information, either interactively or with a script, using the `fc` object for the `stats` command.

For more information about the `fcstat` command, see the `na_fcstat(1)` man page.

Related concepts

[Storage system information and the `stats` command](#) on page 289

Getting SAS adapter and expander information

You can display information about the SAS adapters and expanders used by the storage subsystem.

About this task

You use the `sasstat` or the `sasadmin` command to display information about the SAS adapters and expanders. The `sasstat` command is an alias for the `sasadmin` command.

Step

1. To display information about SAS adapters and expanders, enter one of the following commands:

Command	Description
<code>sasstat expander</code>	Displays configuration information for a SAS expander.
<code>sasstat expander_map</code>	Displays product information for the SAS expanders attached to the SAS channels in the storage system.
<code>sasstat expander_phy_state</code>	Displays the physical state of the SAS expander.
<code>sasstat adapter_state</code>	Displays the state of a logical adapter.
<code>sasstat dev_stats</code>	Displays statistics for the disk drives connected to the SAS channels in the controller.
<code>sasstat shelf</code>	Displays a pictorial representation of the drive population of a shelf.
<code>sasstat shelf_short</code>	Displays the short form of the <code>sasstat shelf</code> command output.

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

Storage system information and 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.

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.

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?
- 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.
- 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.

- 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.
- 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.
- 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.

Next topics

[Viewing the list of available counters](#) on page 290

[Getting detailed information about a counter](#) on page 291

[Using the stats command interactively in singleton mode](#) on page 292

[Using the stats command interactively in repeat mode](#) on page 293

[Collecting system information by using the stats command in background mode](#) on page 294

[Changing the output of a stats command](#) on page 295

[About the stats preset files](#) on page 297

Viewing the list of available counters

You can display the list of counters for a particular object on the command line.

Step

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.

The list of counters is displayed.

```
toaster> 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 detailed information about a counter

Getting detailed information about a counter helps you understand and process the information you get from a `stats` command.

Step

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.

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 of `stats explain counters` command

```

toaster> 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

Using the `stats` command in singleton mode enables you to see a set of information about the system's current state at the command line.

Step

1. Enter the following command:

```
stats show [-e] object_def [object_def...]
```

object_def is one of the following values:

- An object name (*object_name*). For 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*). For 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*). For 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.

The `-e` option allows extended regular expressions (regex) for instance and counter names. With the `-e` option, the instance and counter names are independently interpreted as regular expressions. The asterisk (*) character is still a wildcard representing all instances or counter names. The regular expression is not anchored. You can use `^` to indicate the start of an instance or counter name, and `$` to indicate the end of an instance or counter name.

Examples of stats show command in singleton mode

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

```
toaster> 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

```

The following command returns any counters in the system object ending in "latency".

```

toaster> stats show -e system::latency$
system:system:sys_read_latency:0ms
system:system:sys_write_latency:0ms
system:system:sys_avg_latency:0ms

```

Using the stats command interactively in repeat mode

Using the `stats` command in repeat mode enables you to see a statistic every few seconds.

Step

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 specifies the interval between the iterations of the `stats` command. The default value is one second.

object_def is one of the following values:

- An object name (*object_name*). For 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*). For 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*). For 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 of `stats show` command in repeat mode

The following command shows 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
...	

Related tasks

Using the stats command interactively in singleton mode on page 292

Collecting system information by using the stats command in background mode

You can collect system information from a specified set of counters over time in the background.

About this task

The `stats start` and `stats stop` commands enable you to collect information from a specified set of counters over time in the background. 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, Data ONTAP does not allow you to start more than 50 background `stats` commands, to keep `stats` commands from consuming too many system resources. If you already have 50 background `stats` commands running, you must stop at least one before you can start more. To stop all currently running `stats` commands, you can use the `stats stop -a` command.

See the `na_stats_preset(5)` man page for a list of options.

Steps

1. Enter the following command to start collecting system information:

```
stats start [-I identifier] object_def [object_def...]
```

If you are running only one background `stats` command, you can omit the `-I` parameter.

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 the name of the object.

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. Enter the following command to stop data collection and display the final results:

```
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.

To filter the output of a background `stats` command initiated with a `stats start` command, add `-O name=value` to the `stats stop` command, where *name* is the name of the option you want to omit from the output and the value is `on` or `off`.

Example

The following command filters out all the statistics with zero counter values:

```
stats stop [-I identifier] -O print_zero_values=off
```

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.

Step

1. Do one of the following:

If you want to...	Then...
Send <code>stats</code> output to a file	Add <code>-o filename</code> to your <code>stats show</code> or <code>stats stop</code> command line. <i>filename</i> is the pathname to the file you want to receive the <code>stats</code> output. The file does not need to exist, although any directory in the path must already exist.
Determine whether the output is formatted in rows or columns	Add the <code>-r</code> or <code>-c</code> option to your <code>stats show</code> or <code>stats stop</code> command line. The <code>-r</code> option formats the output in rows and is the default if the <code>-I</code> option is not specified.
Specify a delimiter so that your output can be imported into a database or spreadsheet	Add the <code>-d delimiter</code> option to your <code>stats show</code> or <code>stats stop</code> command line. The <code>-d</code> option only has effect if your output is in column format.
Filter the output of the <code>stats show</code> command	Add <code>-O name=value</code> to the <code>stats show</code> command. <i>name</i> is the name of the option you want to filter and <i>value</i> is <code>on</code> or <code>off</code> .

See the `na_stats_preset(5)` man page for a list of options.

Examples of changing the output of a `stats` command

The following example displays output in rows:

```
toaster> 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.

The following example displays output in columns:

```
toaster> 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.

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
```

The command in the following example filters output of the `stats show` command with zero counter values:

```
stats show -O print_zero_values=off
```

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. 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.

Viewing the list of available presets

The `stats` command supports preset configurations that contain commonly used combinations of statistics and formats.

Step

1. Enter the following command:

```
stats list presets
```

For a description of the preset file format, see the `na_stats_preset(5)` man page.

The list of available presets is displayed.

```
toaster> stats list presets
Stats Presets:
preset1
preset2
other-preset
...
```

How to get system information using perfmon

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.

To use `perfmon` to access storage system performance statistics, you 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.

How to get system information using perfstat

Perfstat is a NetApp tool that reports performance information for both the host and the storage system. It can be run either a UNIX or a Windows host. It collects the performance information and writes it to a text file.

To get more information about `perfstat`, or to download the tool, go to the NOW site and navigate to **Software Downloads > ToolChest**.

Related information

The NOW site - <http://now.netapp.com/>

System performance and resources

Data ONTAP offers features that enable you to manage your system resources, improve your system performance, optimize data layout, and archive performance data.

Next topics

[How to manage storage system resources by using FlexShare](#) on page 299

[How to increase WAFL cache memory](#) on page 307

[Ways to improve storage system performance](#) on page 312

[How to optimize LUN, file, volume, and aggregate layout](#) on page 316

[How to improve read performance](#) on page 330

How to manage storage system resources by using FlexShare

The FlexShare tool is provided by Data ONTAP to enable you to use priorities and hints to increase your control over how your storage system resources are used.

FlexShare uses the following methods:

- Priorities are assigned to volumes, to assign relative priorities between:
 - Different volumes
For example, you could specify that operations on /vol/db are more important than operations on /vol/test.
 - Client data accesses and system operations
For example, you could specify that client accesses are more important than SnapMirror operations.
- Hints are used to affect the way cache buffers are handled for a given volume.

For more information about FlexShare, see the `na_priority(1)` man page.

Next topics

[When to use FlexShare](#) on page 299

[How to use FlexShare](#) on page 302

When to use FlexShare

If your storage system consistently provides the performance required for your environment, then you do not need FlexShare. If, however, your storage system sometimes does not deliver sufficient

performance to some of its users, you can use FlexShare to increase your control over storage system resources to ensure that those resources are being used most effectively for your environment.

The following sample scenarios describe how FlexShare could be used to set priorities for the use of system resources:

- You have different applications on the same storage system. For example, you have a mission-critical database on the same storage system as user home directories. You can use FlexShare to ensure that database accesses are assigned a higher priority than accesses to home directories.
- You want to reduce the impact of system operations (for example, SnapMirror operations) on client data accesses. You can use FlexShare to ensure that client accesses are assigned a higher priority than system operations.
- You have volumes with different caching requirements. For example, if you have a database log volume that does not need to be cached after writing, or a heavily accessed volume that should remain cached as much as possible, you can use the cache buffer policy hint to help Data ONTAP determine how to manage the cache buffers for those volumes.

FlexShare enables you to construct a priority policy that helps Data ONTAP manage system resources optimally for your application environment. FlexShare does not provide any performance guarantees.

Next topics

[FlexShare and priority levels](#) on page 300

[About using FlexShare in storage systems with a high-availability configuration](#) on page 301

[How the default queue works](#) on page 301

[FlexShare and the global io_concurrency option](#) on page 302

Related concepts

[FlexShare and the buffer cache policy values](#) on page 304

Related tasks

[Assigning priority to a volume relative to other volumes](#) on page 302

[Assigning priority to system operations relative to user operations](#) on page 303

FlexShare and priority levels

Priority levels are relative. When you set the priority level of a volume or operation, you are not giving that volume or operation an absolute priority level. Instead, you are providing a hint to Data ONTAP about how to set priorities for accesses to that volume or operations of that type *relative to other accesses or operations*.

For example, setting the priority level of each of your volumes to the highest level will not improve the performance of your system. In fact, doing so would not result in any performance change.

The following table outlines how the listed volume operations affect FlexShare settings.

Volume operation	Effect on FlexShare settings
Deletion	FlexShare settings removed
Rename	FlexShare settings unchanged
FlexClone volume creation	Parent volume settings unchanged FlexShare settings for new FlexClone volume unset (as for a newly created volume)
Copy	Source volume settings unchanged FlexShare settings for destination volume unset (as for a newly created volume)
Offline/online	FlexShare settings preserved

About using FlexShare in storage systems with a high-availability configuration

If you use FlexShare on storage systems with a high-availability configuration, you must ensure that FlexShare is enabled or disabled on *both* nodes. Otherwise, a takeover can cause unexpected results.

After a takeover occurs, the FlexShare priorities you have set for volumes on the node that was taken over are still operational, and the takeover node creates a new priority policy by merging the policies configured on each individual node. For this reason, make sure that the priorities you configure on each node will work well together.

Note: You can use the `partner` command to make changes to FlexShare priorities on a node that has been taken over.

How the default queue works

Understanding how the default priority is used helps you create the optimal priority policy for your storage system.

Any volume that does not have a priority assigned is in the default queue. If you have not assigned a priority to any volume on your system, then all of your volumes are in the default queue, and requests to all volumes are given equal priority.

When you assign a priority to any volume, it is removed from the default queue. Now, requests to that volume are assigned priorities relative to requests for the default queue. But *all of the volumes in the default queue share the resources allocated to the default queue*. So if you assign priorities to a few volumes and leave the rest in the default queue, the results may not be as you expect.

For this reason, once you assign a priority to any volume, you should assign a priority to all volumes whose relative performance you want to control.

For example, you have 30 volumes on your system. You have one volume, `highvol`, that you would like to have faster access to, and one volume, `lowvol`, for which fast access time is not important. You assign a priority of `VeryHigh` to `highvol` and `VeryLow` to `lowvol`. The result of these changes for the `highvol` volume is as expected: when the system is under load, accesses to the `highvol` volume

are given a higher priority than for any other volume. However, accesses to the lowvol volume may still get a higher priority than accesses to the volumes that remain in the default queue (which has a Medium priority). This is because all of the 28 volumes remaining in the default queue are sharing the resources allocated to the default queue.

FlexShare and the global `io_concurrency` option

Disks have a maximum number of concurrent I/O operations they can support; the limit varies according to the disk type. FlexShare limits the number of concurrent I/O operations per volume based on various values including the volume priority and the disk type.

For most customers, the default `io_concurrency` value is correct and should not be changed. If you have nonstandard disks or load, your system performance might be improved by changing the value of the `io_concurrency` option.

For more information about this option, see the `na_priority(1)` man page or contact technical support.

Attention: This option takes effect across the entire system. Use caution when changing its value and monitor system performance to ensure that performance is improved.

How to use FlexShare

You use FlexShare to assign priorities to volume data access, set the volume buffer cache policy, and modify the default priority.

Next topics

[Assigning priority to a volume relative to other volumes](#) on page 302

[Assigning priority to system operations relative to user operations](#) on page 303

[FlexShare and the buffer cache policy values](#) on page 304

[Setting the volume buffer cache policy](#) on page 305

[Removing FlexShare priority from a volume](#) on page 306

[Modifying the default priority](#) on page 306

Assigning priority to a volume relative to other volumes

You can use FlexShare to assign a relative priority to a volume to cause accesses to that volume to receive a priority that is higher or lower than that of other volumes on your storage system.

For best results, when you set the priority of any volume, set the priority of all volumes on the system.

Steps

1. If you have not already done so, ensure that FlexShare is enabled for your storage system by entering the following command:

```
priority on
```

- Specify the priority for the volume by entering the following command:

```
priority set volume vol_name level=priority_level
```

vol_name is the name of the volume for which you want to set the priority.

priority_level is one of the following values:

- VeryHigh
- High
- Medium (default)
- Low
- VeryLow
- A number from 8 (VeryLow) to 92 (VeryHigh)

For more information about the `priority` command, see the `na_priority(1)` man page.

Example

The following command sets the priority level for the `dbvol` volume as high as possible. This causes accesses to the `dbvol` volume to receive a higher priority than accesses to volumes with a lower priority.

```
priority set volume dbvol level=VeryHigh system=30
```

Note: Setting the priority of system operations to 30 does not mean that 30 percent of storage system resources are devoted to system operations. Rather, when both user and system operations are requested, the system operations are selected over the user operations 30 percent of the time, and the other 70 percent of the time the user operation is selected.

- You can optionally verify the priority level of the volume by entering the following command:

```
priority show volume [-v] vol_name
```

Related concepts

[How the default queue works](#) on page 301

Assigning priority to system operations relative to user operations

If system operations (for example, SnapMirror transfers or `ndmptcopy` operations) are negatively affecting the performance of user accesses to the storage system, you can use FlexShare to assign the priority of system operations to be lower than that of user operations for any volume.

Synchronous SnapMirror updates are not considered system operations, because they are performed from NVRAM when the primary operation is initiated. Therefore, synchronous SnapMirror updates are affected by the volume priority of the target volume, but not by the relative priority of system operations for that volume.

Steps

1. If you have not already done so, ensure that FlexShare is enabled for your storage system by entering the following command:

```
priority on
```

2. Specify the priority for system operations for the volume by entering the following command:

```
priority set volume vol_name system=priority_level
```

vol_name is the name of the volume for which you want to set the priority of system operations.

priority_level is one of the following values:

- VeryHigh
- High
- Medium (default)
- Low
- VeryLow
- A number from 4 (VeryLow) to 96 (VeryHigh)

For more information about the `priority` command, see the `na_priority(1)` man page.

Example

The following command sets the priority level for the `dbvol` volume as high as possible while setting system operations for that volume to 30.

```
priority set volume dbvol level=VeryHigh system=30
```

Note: Setting the priority of system operations to 30 does not mean that 30 percent of storage system resources are devoted to system operations. Rather, when both user and system operations are requested, the system operations will be selected over the user operations 30 percent of the time, and the other 70 percent of the time the user operation is selected.

3. You can optionally verify the priority levels of the volume by entering the following command:

```
priority show volume -v vol_name
```

FlexShare and the buffer cache policy values

You can use FlexShare to give Data ONTAP a hint about how to manage the buffer cache for that volume.

Note: This capability only provides a hint to Data ONTAP. Ultimately, Data ONTAP makes the final determination about buffer reuse, based on multiple factors including your input.

The buffer cache policy can be one of the following values:

- keep

This value tells Data ONTAP to wait as long as possible before reusing the cache buffers. This value can improve performance for a volume that is accessed frequently, with a high incidence of multiple accesses to the same cache buffers.

- `reuse`
This value tells Data ONTAP to make buffers from this volume available for reuse quickly. You can use this value for volumes that are written but rarely read, such as database log volumes, or volumes for which the data set is so large that keeping the cache buffers will probably not increase the hit rate.
- `default`
This value tells Data ONTAP to use the default system cache buffer policy for this volume.

Setting the volume buffer cache policy

You can use FlexShare to influence how Data ONTAP determines when to reuse buffers.

Steps

1. If you have not already done so, ensure that FlexShare is enabled for your storage system by entering the following command:

```
priority on
```

2. Specify the cache buffer policy for the volume by entering the following command:

```
priority set volume vol_name cache=policy
```

policy is one of the following policy values:

- `keep`
- `reuse`
- `default`

Example

The following command sets the cache buffer policy for the `testvol1` volume to `keep`, which instructs Data ONTAP not to reuse the buffers for this volume when possible.

```
priority set volume testvol1 cache=keep
```

3. You can optionally verify the priority levels of the volume by entering the following command:

```
priority show volume -v vol_name
```

Related concepts

[FlexShare and the buffer cache policy values](#) on page 304

Removing FlexShare priority from a volume

You can temporarily disable the FlexShare priority for a particular volume, or you can remove the priority completely.

Step

1. Do one of the following:

If you want to...	Then...
Temporarily disable FlexShare priority for a specific volume	Set the service option for that volume to <code>off</code> . Doing so causes that volume to be put back into the default queue.
Completely remove the FlexShare priority settings from a specific volume	Use the <code>priority delete</code> command. Doing so causes that volume to be put back into the default queue.

Example

The following command temporarily disables FlexShare priority for the `testvol1` volume:

```
priority set volume testvol1 service=off
```

Example

The following command completely removes the FlexShare priority settings for the `testvol1` volume:

```
priority delete volume testvol1
```

Modifying the default priority

If you have not assigned a priority to a volume, then that volume is given the default priority for your storage system. The default value for the default priority is Medium. You can change the value of the default priority.

The default priority is also used for all aggregate operations. Changing the default priority to be very high or very low may have unintended consequences.

Step

1. Specify the default volume priority by entering the following command:

```
priority set default option=value [option=value]
```

option is either `level` or `system`, and the possible values for these options are the same as for assigning priorities for a specific volume.

Example

The following command sets the default priority level for volumes to Medium, while setting the default system operations priority to Low.

```
priority set default level=Medium system=Low
```

How to increase WAFL cache memory

You can increase Write Anywhere File Layout (WAFL) cache memory in storage systems running Data ONTAP by using WAFL extended cache and the Performance Acceleration Module (PAM) family. WAFL extended cache is a software component of Data ONTAP and requires a license.

WAFL extended cache provides extra WAFL cache memory to improve the performance of the storage system by reducing the number of disk reads. The `options flexscale` commands enable you to control how user data blocks are cached.

Note: Not all systems have the Performance Acceleration Module installed. Therefore, not all systems can utilize the WAFL extended cache functionality.

WAFL extended cache does not cache data that is stored on an SSD aggregate.

If you use WAFL extended cache on storage systems with a high-availability configuration, you must ensure that the WAFL extended cache options are the same on both nodes. Otherwise, a takeover can result in lower performance due to the lack of WAFL extended cache on the remaining node.

Besides the Data ONTAP options that you can use to manage WAFL extended cache, a diagnostic command is available for sanitizing the Performance Acceleration Module. For more information, see the *Diagnostics Guide*.

Next topics

[Enabling and disabling WAFL extended cache](#) on page 308

[Caching normal user data blocks](#) on page 308

[Caching low-priority user data blocks](#) on page 308

[Caching only system metadata](#) on page 309

[Integrating FlexShare buffer cache policies with WAFL extended cache options](#) on page 310

[Displaying the WAFL extended cache configuration](#) on page 311

[Displaying usage and access information for WAFL extended cache](#) on page 311

Related concepts

[Ways to manage licenses](#) on page 154

Enabling and disabling WAFL extended cache

You can enable the WAFL extended cache functionality for a storage system that has the Performance Acceleration Module installed.

About this task

The `flexscale.enable` option enables or disables the WAFL extended cache functionality. If your storage system does not have the Performance Acceleration Module, the `flexscale.enable` option enables or disables the Predictive Cache Statistics (PCS).

WAFL extended cache needs to be independently licensed. PCS does not require a license.

Step

1. Enter the following command:

```
options flexscale.enable {on|off}
```

The default value is `off`.

Caching normal user data blocks

If you cache normal user data blocks, the WAFL extended cache interprets this setting as the buffer cache policy of `keep` and saves normal user data blocks in the extended cache.

Step

1. To enable or disable caching for normal user data blocks, enter the following command:

```
options flexscale.normal_data_blocks {on|off}
```

The default value is `on`.

When the `flexscale.normal_data_blocks` option is set to `on`, the WAFL extended cache interprets this setting as the buffer cache policy of `keep` and saves normal user data blocks in the extended cache.

If this option is set to `off`, only metadata blocks are cached, except for volumes that have a FlexShare buffer cache policy of `keep`.

Related concepts

[FlexShare and the buffer cache policy values](#) on page 304

Caching low-priority user data blocks

You can cache low-priority user data blocks that are not normally stored by WAFL extended cache. Low-priority blocks include blocks read in large sequential scans that are not normally reused, and blocks that have been written to the storage system through a network-attached storage (NAS) protocol such as Network File System (NFS). Caching low-priority user data blocks is useful if you

have workloads that fit within WAFL extended cache memory and if the workloads consist of either write followed by read or large sequential reads.

About this task

You can cache low-priority user data blocks (setting `flexscale.lopri_blocks` to `on`) only if you also cache normal user data blocks (by setting `flexscale.normal_data_blocks` to `on`).

Step

1. To control whether low-priority user data blocks are cached, enter the following command:

```
options flexscale.lopri_blocks {on|off}
```

The default value is `off`.

Setting the option to `on` caches low-priority user data blocks.

Related tasks

[Caching normal user data blocks](#) on page 308

Caching only system metadata

If the working set of the storage system is very large, such as a large e-mail server, you can cache only system metadata in WAFL extended cache memory by turning off both normal user data block caching and low-priority user data block caching.

About this task

When you cache only system metadata, with both `flexscale.normal_data_blocks` and `flexscale.lopri_blocks` set to `off`, WAFL extended cache interprets this setting as the buffer cache policy of `reuse` and does not save normal data blocks or low-priority blocks in the extended cache.

Steps

1. Enter the following command to turn off normal user data block caching:
2. Enter the following command to turn off low-priority user data block caching:

```
options flexscale.normal_data_blocks off
```

```
options flexscale.lopri_blocks off
```

Related concepts

[FlexShare and the buffer cache policy values](#) on page 304

Related tasks

[Setting the volume buffer cache policy](#) on page 305

Integrating FlexShare buffer cache policies with WAFL extended cache options

For additional cache control, you can integrate FlexShare buffer cache policies with the WAFL extended cache options (`flexscale.normal_data_blocks` and `flexscale.lopri_blocks`). Doing so allows you to set caching policies on specific volumes. You can choose to enable only the FlexShare buffer cache policies without enabling all other FlexShare options.

Steps

1. If you have not already done so, ensure that FlexShare is enabled for your storage system by entering the following command:

```
priority on
```

2. To enable only the FlexShare buffer cache policies and not other FlexShare options, enter the following command:

```
priority set enabled_components=cache
```

3. To specify the caching policy for a specific volume, enter the following command:

```
priority set volume myvol cache=policy
```

policy can be `keep` or `reuse`.

When you cache metadata for the system, setting `cache` to `keep` for *myvol* enables you to cache normal user data for only *myvol*.

Note: If you cache normal or low-priority user data for the system, setting `cache` to `keep` for *myvol* has no effect on the specified volume.

When you cache normal or low-priority user data for the system, setting `cache` to `reuse` for *myvol* enables you to cache metadata for only *myvol*.

Examples of integrating FlexShare policies with the WAFL extended cache options

After the FlexShare buffer cache policies are enabled, the following commands enable you to cache metadata for the system as a whole but cache normal user data for the *myvol* volume:

```
options flexscale.normal_data_blocks off
```

```
options flexscale.lopri_blocks off
```

```
priority set volume myvol cache=keep
```

After the FlexShare buffer cache policies are enabled, the following commands enable you to cache normal user data for the system as a whole but cache only metadata for *myvol*:

```
options flexscale.normal_data_blocks on
```

```
options flexscale.lopri_blocks off
priority set volume myvol cache=reuse
```

Related tasks

[Setting the volume buffer cache policy](#) on page 305

Displaying the WAFL extended cache configuration

Data ONTAP enables you to display configuration information for WAFL extended cache.

Step

1. Enter the following command:

```
stats show -p flexscale
```

Displaying usage and access information for WAFL extended cache

You can display usage and access information for WAFL extended cache, have output produced periodically, and terminate the output after a specified number of iterations.

Step

1. Enter the following command:

```
stats show -p flexscale-access [-i interval] [-n num]
```

- If no options are used, a single one-second snapshot of statistics is used.
- `-i interval` specifies that output is to be produced periodically, with an interval of `interval` seconds between each set of output.
- `-n num` terminates the output after `num` number of iterations, when the `-i` option is also used. If no `num` value is specified, the output runs forever until a user issues a break.
- Press Ctrl-C to interrupt output.

Example

The following example shows sample output from the `stats show -p flexscale-access` command:

Cache	Reads				Writes				Disk				
Read	Usage	Hit	Meta	Miss	Hit	Evict	Inval	Insrt	Chain	Blcks	Chain	Blcks	Replcd
%	/s	/s	/s	%	/s	/s	/s	/s	/s	/s	/s	/s	/s
0	581	0	83	87	0	604	13961	579	581	218	13960	552	
0	777	0	133	85	0	121	21500	773	777	335	21494	744	

0	842	0	81	91	0	1105	23844	837	842	372	23845	812
0	989	0	122	89	0	0	23175	981	989	362	23175	960

Example

The following command displays access and usage information for WAFL extended cache once every 10 seconds for 5 times:

```
stats show -p flexscale-access -i 10 -n 5
```

Ways to improve storage system performance

You can take configuration procedures to improve your system's performance.

Next topics

[About balancing NFS traffic on network interfaces](#) on page 312

[How to ensure reliable NFS traffic by using TCP](#) on page 312

[Avoiding access time update for inodes](#) on page 313

[Adding disks to a disk-bound aggregate](#) on page 313

[About sizing aggregates appropriately](#) on page 314

[About putting cards into the correct slots](#) on page 314

[Maintaining adequate free blocks and free inodes](#) on page 314

[About optimizing LUN, file, and volume layout](#) on page 315

[Using oplocks for CIFS storage systems](#) on page 315

[Increasing the TCP window size for CIFS or NFS](#) on page 315

[About backing up by using qtrees](#) on page 316

About balancing NFS traffic on network interfaces

You can 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.

How to ensure reliable NFS traffic by using TCP

With faster NICs and switches, you are advised to support NFSv2 or NFSv3 protocol over TCP rather than over UDP. NFSv4 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.

About this task

Consider setting the `no_atime_update` option to `on` to prevent updates 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.

Step

1. Enter the following command:

```
vol options volname no_atime_update on
```

Adding disks to a disk-bound aggregate

If you have a single traditional volume or single-aggregate storage system, you can determine the fraction of time that the busiest disk is active and add disks to the traditional volume or aggregate if necessary.

About this task

Consider taking advantage of a 64-bit aggregate, which can hold more disks than a 32-bit aggregate. For information on 64-bit aggregates, see the *Data ONTAP 7-Mode Storage Management Guide*.

Steps

1. Enter the following command to determine the fraction of time that the busiest disk is active:

```
sysstat -u
```

2. If the fraction is greater than 80 percent, add disks to the traditional volume or aggregate by entering the following command:

```
aggr add aggrname disk-list
```

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

About 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).

About 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.

For information about card placement, see the *System Configuration Guide*.

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. You should check to ensure that you system has adequate free blocks and free inodes.

Steps

1. Enter one of the following commands:

If you want to check ...	Enter this command...
Free blocks	<code>df</code>
Free inodes	<code>df -i</code>

2. Do one of the following as necessary:

- 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.

For more information about deleting Snapshot copies, see the `na_snap(1)` man page and the *Data ONTAP 7-Mode Block Access Management Guide for iSCSI and FC*.

For more information about the `maxfiles` command, see the `na_maxfiles(1)` man page.

About optimizing LUN, file, and volume layout

If read performance on a particular large file or LUN degrades over time, consider using 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.

Related concepts

[How to optimize LUN, file, volume, and aggregate layout](#) on page 316

[What a reallocation scan is](#) on page 317

Using oplocks for CIFS storage systems

Oplocks (opportunistic locks) allow CIFS clients to cache more data locally, reducing traffic to the storage system.

Step

1. If your system is running CIFS and is not in a database environment, enter the following command to set oplocks:

```
options cifs.oplocks.enable on
```

Attention: If your system is in a database environment, ensure that the oplocks are *not* set.

For more information about the `cifs.oplocks.enable` option, see the `na_options(1)` man page.

Increasing the TCP window size for CIFS or NFS

The TCP window size controls the number of TCP messages that can be transmitted between the storage system and the client at one time. Increasing the TCP receive window size to its maximum setting on both the system and the client can improve performance for large transfers, provided that packet loss is not taking place and the client's send window is large.

About this task

You should call technical support before changing this value.

Steps

1. Do one of the following:

To maximize the TCP window size on a storage system running...	Enter this command...
---	------------------------------

CIFS	<code>options cifs.tcp_window_size 8388608</code>
------	---

To maximize the TCP window size on a storage system running... **Enter this command...**

NFS `options nfs.tcp.recvwindowsize 8388608`

- For the CIFS protocol, the default is 17,520 bytes. The number of bytes must be between 1,600 and 8,388,608.
- For the NFS protocol, the default is 65,940 bytes. The number of bytes must be between 8,760 and 8,388,608.

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.

2. 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).

About 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.

For more information about the `dump` command, see the `na_dump(1)` man page.

How to optimize LUN, file, volume, and aggregate layout

You can optimize the existing layout of a LUN, a file, a volume, or an aggregate.

Optimizing the existing layout of a LUN, file, or volume improves the sequential read performance of host applications that access data on the storage system. Write performance may also be improved as a result of file reallocation. Optimizing the layout of a volume is equivalent to optimizing all files and LUNs in the volume.

Optimizing the existing layout of an aggregate improves contiguous free space in the aggregate, hence improving the layout, and usually the performance, of future writes to volumes in the aggregate. Optimizing the aggregate layout is not equivalent to optimizing all the volumes in the aggregate.

Note: "LUNs" in this context refers to the LUNs that Data ONTAP serves to clients, not to the array LUNs used for storage on a storage array.

Next topics

What a reallocation scan is on page 317

Reasons to use LUN, file, or volume reallocation scans on page 318

Reasons to use aggregate reallocation scans on page 318

Reasons to use physical reallocation scans on page 318

How a reallocation scan works on page 319

How you manage reallocation scans on page 320

How to use reallocation scans most efficiently on page 329

What a reallocation scan is

A reallocation scan evaluates how the blocks are laid out on disk in a LUN, file, volume, or aggregate, and rearranges them if necessary.

Data ONTAP performs the scan as a background task, so applications can rewrite blocks in the LUN, file, volume, or aggregate during the scan. Repeated layout checks during a file, LUN, or volume reallocation scan ensure that the sequential block layout is maintained during the current scan.

A reallocation scan does not necessarily rewrite every block in the LUN, file, or volume. Rather, it rewrites whatever is required to optimize the block layout.

A file reallocation scan using `reallocate start` or `reallocate start -p` does not rearrange blocks that are shared between files by deduplication on deduplicated volumes. Because a file reallocation scan does not predictably improve read performance when used on deduplicated volumes, it is best not to perform file reallocation on deduplicated volumes. If you want your files to benefit from a reallocation scan, store them on volumes that are not enabled for deduplication.

Note: Output of a reallocation scan goes to the system log. You can view the current status by using the `reallocate status` command.

The following general recommendations apply to a file, volume, or aggregate reallocation scan:

- The best time to run a reallocation scan is when the storage system is relatively idle or when minimal write requests are going to the target volume.
- Reallocation scans will not run if there is less than 10 percent free space (excluding the Snapshot reserve) in the active file system on the target volume or aggregate. The more free space the target has, the more effective the reallocation scan is.
- Check to make sure that the target volume's space guarantee is enabled so that the reallocation scan does not cause an overcommitment of the volume's storage space. For information about space guarantees, see the *Data ONTAP 7-Mode Storage Management Guide*.
- Before a reallocation scan, minimize the number of Snapshot copies in the target volume or aggregate by deleting unwanted Snapshot copies.

When you use `reallocate start` without the `-p` option, a reallocation scan duplicates blocks that are held in a Snapshot copy, so a file might use more space after the scan. When you use `reallocate start` with the `-p` option, blocks are moved, and the file takes up less additional space after the scan.

- If a volume you want to reallocate involves SnapMirror, reallocate the source volume instead of the destination volume.

Related concepts

[Reasons to use physical reallocation scans](#) on page 318

[How you manage reallocation scans](#) on page 320

Reasons to use LUN, file, or volume reallocation scans

You run LUN, file, or volume reallocation scans to ensure that blocks in a LUN, file, or volume are laid out sequentially.

If a LUN, file, or volume is not laid out in sequential blocks, sequential read commands take longer to complete because each command might require an additional disk seek operation. Sequential block layout may improve the sequential read performance, and usually the write performance, of host applications that access data on the storage system.

You run a LUN, file, or volume reallocation using the `reallocate start` command. If you add disks to an aggregate, you can redistribute the data equally across all of the disks in the aggregate using the `reallocate start -f` command.

Note: A volume reallocation scan computes the average level of layout optimization over all the files in the volume. Therefore, a volume reallocation works best if a volume has many files or LUNs with similar layout characteristics.

Reasons to use aggregate reallocation scans

You run aggregate reallocation scans to optimize the location of physical blocks in the aggregate. Doing so increases contiguous free space in the aggregate.

You run an aggregate reallocation scan using the `reallocate start -A` command.

Aggregate reallocation does not optimize the existing layout of individual files or LUNs. Instead, it optimizes the free space where future blocks can be written in the aggregate. Therefore, if the existing layout for a file, LUN, or volume is not optimal, run a file, LUN, or volume reallocation scan. For instance, after adding new disks to an aggregate, if you want to ensure that blocks are laid out sequentially throughout the aggregate, you should use `reallocate start -f` on each volume instead of `reallocate start -A` on the aggregate.

Note: Aggregate reallocation is not supported on aggregates created by versions of Data ONTAP earlier than 7.2. If you try to perform an aggregate reallocation on such an aggregate, you receive a message saying that the reallocation is not supported. For more information, see the `na_reallocate(1)` man page.

Reasons to use physical reallocation scans

A physical reallocation (using the `-p` option of the `reallocate start` command) reallocates user data on the physical blocks in the aggregate while preserving the logical block locations within a

FlexVol volume. You can perform physical reallocation with FlexVol volumes or files and LUNs within FlexVol volumes.

Physical reallocation might reduce the extra storage requirements in a FlexVol volume when reallocation is run on a volume with Snapshot copies. It might also reduce the amount of data that needs to be transmitted by SnapMirror on its next update after reallocation is performed on a SnapMirror source volume.

Physical reallocation is not supported on FlexVol volumes or on files and LUNs within FlexVol volumes that are in an aggregate created by a version of Data ONTAP earlier than version 7.2.

Physical reallocation is also not supported on RAID0.

Note: Using the `-p` option might cause a performance degradation when reading older Snapshot copies, if the volume has significantly changed after reallocation. Performance might be impacted when reading files in the `.snapshot` directory, accessing a LUN backed up by a Snapshot copy, or reading a qtree SnapMirror (QSM) destination. This performance degradation does not occur with whole-volume reallocation.

How a reallocation scan works

Data ONTAP performs file reallocation scans and aggregate reallocation scans in different ways.

- Data ONTAP performs a file reallocation scan as follows:
 1. Scans the current block layout of the LUN.
 2. Determines the level of optimization of the current layout on a scale of 3 (moderately optimal) to 10 (not optimal).
 3. Performs one of the following tasks, depending on the optimization level of the current block layout:
 - If the layout is optimal, the scan stops.
 - If the layout is not optimal, blocks are reallocated sequentially.

Note: In addition to the specified threshold level, Data ONTAP also includes “hot spots” in its calculation of whether to start a file reallocation. As a result, Data ONTAP might start a reallocation when the average optimization is better than the threshold but a small percentage of the total data is very poorly optimized.
 4. Scans the new block layout.
 5. Repeats steps 2 and 3 until the layout is optimal.
- Data ONTAP performs an aggregate reallocation scan by scanning through an aggregate and reallocating blocks as necessary to improve free-space characteristics.

The rate at which the reallocation scan runs (the blocks reallocated per second) depends on CPU and disk loads. For example, if you have a high CPU load, the reallocation scan will run at a slower rate, so as not to impact system performance.

How you manage reallocation scans

To manage reallocation scans, you must enable reallocation scans on your storage system. Then you define a reallocation scan to run at specified intervals or on a specified schedule.

You manage reallocation scans by performing the following tasks:

- First, enable reallocation scans.
- Then, either define a reallocation scan to run at specified intervals (such as every 24 hours), or define a reallocation scan to run on a specified schedule that you create (such as every Thursday at 3:00 p.m.).

You can define only one reallocation scan per file, LUN, volume, or aggregate. You can, however, define reallocation scans for both the aggregate (to optimize free space layout) and the volumes in the same aggregate (to optimize data layout).

You can also initiate scans at any time, force Data ONTAP to reallocate blocks sequentially regardless of the optimization level of the LUN layout, and monitor and control the progress of scans.

A file or LUN reallocation scan is not automatically deleted when you delete its corresponding file or LUN. This allows you to reconstruct the file or LUN without having to recreate its reallocation scan. If the file or LUN has not been recreated in time for the next scheduled run of the reallocation scan, the storage system console displays an error message. A volume or aggregate reallocation scan is automatically deleted when you delete its corresponding volume or aggregate.

You can perform reallocation scans on LUNs or aggregates when they are online. You do not have to take them offline. You also do not have to perform any host-side procedures when you perform reallocation scans.

Next topics

[Enabling reallocation scans](#) on page 321

[Defining a LUN, file, or volume reallocation scan](#) on page 321

[Defining an aggregate reallocation scan](#) on page 322

[Creating a reallocation scan schedule](#) on page 323

[Deleting a reallocation scan schedule](#) on page 324

[Starting a one-time reallocation scan](#) on page 324

[Performing a full reallocation scan of a LUN, file, or volume](#) on page 325

[Performing a measure-only reallocation scan of a LUN or volume](#) on page 326

[Quiescing a reallocation scan](#) on page 327

[Restarting a reallocation scan](#) on page 327

[Displaying the status of a scan](#) on page 328

[Deleting a reallocation scan](#) on page 328

[Disabling reallocation scans](#) on page 329

Enabling reallocation scans

Reallocation scans are disabled by default. You must enable reallocation scans globally on the storage system before you run a scan or schedule regular scans.

Step

1. On the storage system's command line, enter the following command:

```
reallocate on
```

Defining a LUN, file, or volume reallocation scan

After reallocation is enabled on your storage system, you define a reallocation scan for the LUN, file, or volume on which you want to perform a reallocation scan.

Step

1. On the storage system's command line, enter the following command:

```
reallocate start [-t threshold] [-n] [-p] [-i interval] pathname
```

- *-t threshold* is a number between 3 (layout is moderately optimal) and 10 (layout is not optimal). The default is 4.

A scan checks the block layout of a LUN, file, or volume before reallocating blocks. If the current layout is below the threshold, the scan does not reallocate blocks in the LUN, file, or volume. If the current layout is equal to or above the threshold, the scan reallocates blocks in the LUN, file, or volume.

Note: Because Data ONTAP also includes “hot spots” in its calculation of whether to start a LUN, file, or volume reallocation, the system might start a reallocation when the average optimization is better than the threshold but a small percentage of the total data is very poorly optimized.

- *-n* reallocates blocks in the LUN, file, or volume without checking its layout.
- *-p* reallocates user data on the physical blocks in the aggregate while preserving the logical block locations within a FlexVol volume. You can use this option only with FlexVol volumes or with files and LUNs within FlexVol volumes.
- *-i interval* is the interval, in hours, minutes, or days, at which the scan is performed. The default interval is 24 hours. You specify the interval as follows:

```
[m | h | d]
```

For example, **30m** is a 30-minute interval.

The countdown to the next scan begins only after the first scan is complete. For example, if the interval is 24 hours and a scan starts at midnight and lasts for an hour, the next scan begins at 1:00 a.m. the next day—24 hours after the first scan is completed.

- *pathname* is the path to the LUN, file, or volume on which you want to perform a reallocation scan.

Example

The following commands create a new LUN and a normal reallocation scan that runs every 24 hours.

```
lun create -s 100g /vol/vol2/lun0
reallocate start /vol/vol2/lun0
```

Related concepts

[How you manage reallocation scans](#) on page 320

Related tasks

[Creating a reallocation scan schedule](#) on page 323

[Enabling reallocation scans](#) on page 321

Defining an aggregate reallocation scan

If reallocation has been enabled on your storage system, you can initiate an aggregate reallocation scan to optimize the location of physical blocks in the aggregate, thus increasing contiguous free space in the aggregate.

An aggregate reallocation scan is reallocation of free space and is not equivalent to file reallocation. In particular, after adding new disks to an aggregate, if you want to ensure that blocks are laid out sequentially throughout the aggregate, you should use `reallocate start -f` on each volume instead of `reallocate start -A` on the aggregate.

Because blocks in an aggregate Snapshot copy will not be reallocated, consider deleting aggregate Snapshot copies before performing aggregate reallocation to allow the reallocation to perform better.

Volumes in an aggregate on which aggregate reallocation has started but has not successfully completed will have the `active_redirect` status. Read performance of such volumes may be degraded until aggregate reallocation has successfully completed. Volumes in an aggregate that has previously undergone aggregate reallocation have the `redirect` status. For more information, see the `na_vol(1)` man page.

Step

1. On the storage system's command line, enter the following command:

```
reallocate start -A [-i interval] aggr_name
```

- `-i interval` is the interval, in hours, minutes, or days, at which the scan is performed. The default interval is 24 hours. You specify the interval as follows:

```
[m | h | d]
```

For example, `30m` is a 30-minute interval.

The countdown to the next scan begins only after the first scan is complete. For example, if the interval is 24 hours and a scan starts at midnight and lasts for an hour, the next scan begins at 1:00 a.m. the next day—24 hours after the first scan is completed.

- *aggr_name* is the name of the aggregate on which you want to perform a reallocation scan.

Example

The following example initiates an aggregate reallocation scan that runs every 24 hours.

```
reallocate start -A my_aggr
```

Related concepts

[Reasons to use aggregate reallocation scans](#) on page 318

Related tasks

[Performing a full reallocation scan of a LUN, file, or volume](#) on page 325

[Creating a reallocation scan schedule](#) on page 323

Creating a reallocation scan schedule

You can run reallocation scans according to a schedule. The schedule you create replaces any interval you specified when you entered the `reallocate start` command or the `reallocate start -A` command.

If the reallocation scan job does not already exist, use `reallocate start` first to define the reallocation scan.

Step

1. Enter the following command:

```
reallocate schedule [-s schedule] pathname | aggr_name
```

-s schedule is a string with the following fields:

```
minute hour day_of_month day_of_week
```

- *minute* is a value from 0 to 59.
- *hour* is a value from 0 (midnight) to 23 (11:00 p.m.).
- *day_of_month* is a value from 1 to 31.
- *day_of_week* is a value from 0 (Sunday) to 6 (Saturday).

A wildcard character (*) indicates every value for that field. For example, a * in the *day_of_month* field means every day of the month. You cannot use the wildcard character in the *minute* field.

You can enter a number, a range, or a comma-separated list of values for a field. For example, entering “0,1” in the *day_of_week* field means Sundays and Mondays. You can also define a range of values. For example, “0-3” in the *day_of_week* field means Sunday through Wednesday.

pathname is the path to the LUN, file, or volume for which you want to create a reallocation scan schedule.

aggr_name is the name of the aggregate for which you want to create a reallocation scan schedule.

Example

The following example schedules a LUN reallocation scan for every Saturday at 11:00 PM.

```
reallocate schedule -s "0 23 * 6" /vol/myvol/lun1
```

Deleting a reallocation scan schedule

You can delete an existing reallocation scan schedule that is defined for a LUN, a file, a volume, or an aggregate. If you delete a schedule, the scan runs according to the interval that you specified when you initially defined the scan using the `reallocate start` command or the `reallocate start -A` command.

A file or LUN reallocation scan is not automatically deleted when you delete its corresponding file or a LUN. A volume or aggregate reallocation scan is automatically deleted when you delete its corresponding volume or aggregate.

Step

1. Enter the following command:

```
reallocate schedule -d pathname | aggr_name
```

pathname is the path to the LUN, file, or volume on which you want to delete a reallocation scan schedule.

aggr_name is the name of the aggregate on which you want to delete a reallocation scan schedule.

Example

```
reallocate schedule -d /vol/myvol/lun1
```

```
reallocate schedule -d my_aggr
```

Starting a one-time reallocation scan

You can perform a one-time reallocation scan on a LUN, a file, a volume, or an aggregate. This type of scan is useful if you do not want to schedule regular scans for a particular LUN, file, volume, or aggregate.

Step

1. Enter one of the following commands:

To perform a one-time reallocation scan on ...	Enter ...
a LUN, file, or volume	<code>reallocate start -o -n <i>pathname</i></code>
an aggregate	<code>reallocate start -A -o <i>aggr_name</i></code>

- `-o` performs the scan only once.
- `-n` performs the scan without checking the layout of the LUN, file, or volume.

Example

The following example initiates a one-time reallocation scan on the `my_aggr` aggregate.

```
reallocate start -A -o my_aggr
```

Performing a full reallocation scan of a LUN, file, or volume

You can perform a scan that reallocates every block in a LUN, file, or volume regardless of the current layout by using the `-f` option of the `reallocate start` command. A full reallocation optimizes layout more aggressively than a normal reallocation scan. A normal reallocation scan moves blocks only if the move improves the layout of a LUN, file, or volume. A full reallocation scan always moves blocks, unless the move makes the layout even worse.

Using the `-f` option of the `reallocate start` command implies the `-o` and `-n` options. This means that the full reallocation scan is performed only once, without checking the layout first.

You might want to perform this type of scan if you add a new RAID group to a volume and you want to ensure that blocks are laid out sequentially throughout the volume or LUN.

Attention: You cannot perform a full reallocation (using the `-f` option) on an entire volume that has existing Snapshot copies, unless you also perform a physical reallocation (using the `-p` option). Otherwise, an error message is displayed. If you do a full reallocation on a file or LUN without the `-p` option, you might end up using significantly more space in the volume, because the old, unoptimized blocks are still present in the Snapshot copy after the scan. For individual LUNs or files, avoid transferring large amounts of data from the Snapshot copy to the active file system unless absolutely necessary. The greater the differences between the LUN or file and the Snapshot copy, the more likely the full reallocation will be successful.

If a full reallocation scan fails because of space issues, consider performing reallocation scans on a per-file basis, by using `reallocate start file_pathname` without any options. However, if the space issue is caused by a full reallocation on a file or LUN that was performed without the `-p` option, a long-term solution is to wait until the Snapshot rotation has freed space on the volume and then to rerun the full reallocation scan with the `-p` option.

Step

1. Enter the following command:

```
reallocate start -f [-p] pathname | vol/volname
```

-p reallocates user data on the physical blocks in the aggregate while preserving the logical block locations within a FlexVol volume. You can use this option only with FlexVol volumes or with files and LUNs within FlexVol volumes.

Performing a measure-only reallocation scan of a LUN or volume

A measure-only reallocation scan is similar to a normal reallocation scan except that only the check phase is performed. It allows the optimization of the LUN, file, or volume to be tracked over time or measured ad-hoc.

A measure-only reallocation scan checks the layout of a LUN, file, or volume. If the layout measurement becomes less optimal than the threshold (specified by the `-t threshold` option), or if a portion of the data is very poorly optimized, the log message advises you to consider performing a LUN, file, or volume reallocation (using the `reallocate start` command) to optimize the layout.

For scheduled measure-only reallocation scans, the optimization of the last completed check is saved and may be viewed at any time by using `reallocate status`.

Additional detailed information about the layout of the LUN, file, or volume is logged if you use the `-l logfile` option.

Step

1. Enter the following command:

```
reallocate measure [-l logfile] [-t threshold] [-i interval] [-o]
pathname | /vol/volname
```

- `-l logfile` is the file where information about the layout is recorded. If `logfile` is specified, information about the layout is recorded in the file.
- `-t threshold` is a number between 3 (layout is moderately optimal) and 10 (layout is not optimal). The default is 4. When the layout becomes less optimal than the threshold level, the layout of the LUN, file, or volume is considered unoptimized, and the log message advises you to consider performing a LUN, file, or volume reallocation.

Note: Because Data ONTAP also includes “hot spots” in its calculation of whether to start a reallocation, the log message might advise you to consider performing a reallocation when the average optimization is better than the threshold but a small percentage of the total data is very poorly optimized.

- `-i interval` is the interval, in minutes, hours, or days, at which the scan is performed. A measure-only reallocation scan runs periodically at a system-defined interval, but depending on the system configuration and write/read workload, you can change the job interval with the `-i` option. You specify the interval as follows:

```
[m | h | d]
```

For example, `30m` is a 30-minute interval.

The countdown to the next scan begins only after the first scan is complete. For example, if the interval is 24 hours and a scan starts at midnight and lasts for an hour, the next scan begins at 1:00 a.m. the next day—24 hours after the first scan is completed.

- `-o` performs the scan only once, after which the scan is automatically removed from the system.

Example

The following example measures the optimization of the `dblun` LUN once and records detailed information about the measurement in the `measure_log_dblun` log.

```
reallocate measure -o -l /vol/logs/measure_log_dblun/vol/dbvol/dblun
```

After a measure-only reallocation scan, the optimization information is logged via EMS in the system log files.

Quiescing a reallocation scan

You can quiesce (temporarily stop) a reallocation scan that is in progress and restart it later. A file, LUN, or volume reallocation scan restarts from the beginning of the reallocation process. An aggregate reallocation scan restarts from where it stopped. For example, if you want to back up a LUN or an aggregate but a scan is already in progress, you can quiesce the scan.

Step

1. Enter the following command:

```
reallocate quiesce pathname | aggr_name
```

pathname is the path to the LUN, file, or volume, and *aggr_name* is the name of the aggregate for which you want to quiesce the reallocation scan.

Restarting a reallocation scan

You might need to restart a scan that was previously quiesced or a scheduled scan that is currently idle.

You might restart a scan for the following reasons:

- You quiesced the scan by using the `reallocate quiesce` command, and you want to restart it.
- You have a scheduled scan that is idle (it is not yet time for it to run again), and you want to run it immediately.

Step

1. Enter the following command:

```
reallocate restart [-i] pathname | aggr_name
```

- The `-i` option ignores the checkpoint and starts the job at the beginning.

- *pathname* is the path to the LUN, file, or volume on which you want to restart the reallocation scan.
- *aggr_name* is the name of the aggregate on which you want to restart the reallocation scan.

The command restarts a quiesced scan. If there is a scheduled scan that is idle, the `reallocate restart` command runs the scan.

Displaying the status of a scan

You can display the status of a scan, including the state, schedule, interval, optimization, and log file.

Step

1. Enter the following command:

```
reallocate status [-v] [pathname | aggr_name]
```

- *pathname* is the path to the LUN, file, or volume for which you want to see reallocation scan status.
- *aggr_name* is the name of the aggregate for which you want to see reallocation scan status.
- If you do not specify a value for *pathname* or *aggr_name*, then the status for all scans is displayed.

The `reallocate status` command displays the following information:

- State—whether the scan is in progress or idle.
- Schedule—schedule information about the scan. If there is no schedule, then the `reallocate status` command displays `n/a`.
- Interval—intervals at which the scan runs, if there is no schedule defined.
- Optimization—information about the LUN layout.
- Logfile—the name of the logfile for a measure-only scan, if a detail logfile was specified.
- Hot spot optimization—displayed only for scheduled reallocation jobs.

Deleting a reallocation scan

You can permanently delete a scan you defined for a LUN, a file, a volume, or an aggregate. You can also stop any scan that is in progress on the LUN, file, volume, or aggregate.

Step

1. Enter the following command:

```
reallocate stop pathname | aggr_name
```

pathname is the path to the LUN, file, or volume and *aggr_name* is the name of the aggregate on which you want to delete a scan.

The `reallocate stop` command stops and deletes any scan on the LUN, file, volume, or the aggregate, including a scan in progress, a scheduled scan that is not running, or a scan that is quiesced.

Disabling reallocation scans

You can disable reallocation on the storage system. When you disable reallocation scans, you cannot start or restart any new scans. Any scans that are in progress are stopped.

Step

1. Enter the following command:

```
reallocate off
```

Note: If you want to reenable reallocation scans at a later date, use the `reallocate on` command.

How to use reallocation scans most efficiently

To maximize efficiency, you should follow certain guidelines when using reallocation scans.

The following are good practices to follow when you choose to use the `reallocate` command:

- You should define a reallocation scan when you first create the LUN, file, or volume. This ensures that the layout remains optimized as a result of regular reallocation scans.
- You should define regular reallocation scans by using either intervals or schedules. This ensures that the layout of the LUN, file, or volume remains optimized. If you wait until most of the blocks in the layout of the LUN, file, or volume are not sequential, a reallocation scan will take more time.
- You should define intervals according to the type of read/write activity associated with the LUN, file, or volume:
 - Long intervals—You should define long reallocation scan intervals for LUNs, files, or volumes in which the data changes slowly, for example, when data changes as a result of infrequent large write operations.
 - Short intervals—You should define short reallocation scan intervals for LUNs, files, or volumes that are characterized by workloads with many small random write and many sequential read operations. These types of LUNs, files, or volumes might become heavily fragmented over a shorter period of time.
- If you do not know the type of read/write activity associated with the LUNs, files, or volumes, you can choose to rely on the default layout of the system.

How to improve read performance

There are some tasks you can perform to improve the read performance of your storage system.

Next topics

[About read reallocation](#) on page 330

[About improving Microsoft Exchange read performance](#) on page 331

About read reallocation

For workloads that perform a mixture of random writes and large and multiple sequential reads, read reallocation improves the file's layout and sequential read performance.

When you enable read reallocation, Data ONTAP analyzes the parts of the file that are read sequentially. If the associated blocks are not already largely contiguous, Data ONTAP updates the file's layout by rewriting those blocks to another location on disk. The rewrite improves the file's layout, thus improving the sequential read performance the next time that section of the file is read.

However, read reallocation might result in a higher load on the storage system. Also, unless you set `vol options vol-name read_realloc to space_optimized`, read reallocation might result in more storage use if Snapshot copies are used.

If you want to enable read reallocation but storage space is a concern, you can enable read reallocation on FlexVol volumes by setting `vol options vol-name read_realloc to space_optimized` (instead of `on`). Setting the option to `space_optimized` conserves space but results in degraded read performance through the Snapshot copies. Therefore, if fast read performance through Snapshot copies is a high priority to you, do not use `space_optimized`.

Read reallocation might conflict with deduplication by adding new blocks that were previously consolidated during the deduplication process. A deduplication scan might also consolidate blocks that were previously rearranged by the read reallocation process, thus separating chains of blocks that were sequentially laid out on disk. Therefore, since read reallocation does not predictably improve the file layout and the sequential read performance when used on deduplicated volumes, performing read reallocation on deduplicated volumes is not recommended. Instead, for files to benefit from read reallocation, they should be stored on volumes that are not enabled for deduplication.

The read reallocation function is not supported on FlexCache volumes. If file fragmentation is a concern, enable the read reallocation function on the original server volume.

Enabling and disabling read reallocation

You can enable read reallocation to improve subsequent read performance of a file.

Step

1. Enter the following command:

```
vol options vol-name read_realloc [on | space_optimized | off]
```

- `on` enables read reallocation for the volume to improve its subsequent read performance. Enabling read reallocation might help workloads that perform a mixture of random writes and large and multiple sequential reads. However, enabling read reallocation might increase the number of disk operations performed on the storage system.
- `space_optimized` also enables read reallocation but can be used only on FlexVol volumes. Using `space_optimized` might be useful if the FlexVol volume has Snapshot copies or is a SnapMirror source. When you use `space_optimized`, the extent update does not result in duplicated Snapshot blocks in the active file system, thus conserving space in the volume. Also, `space_optimized` might reduce the amount of data that SnapMirror needs to move on the next update. However, `space_optimized` might result in degraded Snapshot read performance.
`space_optimized` is not supported if `vol-name` is in an aggregate that was either created prior to Data ONTAP 7.2 or once reverted to a version earlier than Data ONTAP 7.2.
- `off` disables read reallocation for the volume. By default, read reallocation is disabled.

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

About improving Microsoft Exchange read performance

In Microsoft Exchange environments, you can use the Exchange `eseutil` tool to perform database scans for validation purposes. Exchange database scans usually access data by using a sequential read pattern. By enabling extents, you improve Exchange sequential read performance and database validation time.

An extent is a group of user-level data blocks that are aligned and contiguous. When you enable extents, Data ONTAP processes write operations by creating groups of contiguous data blocks on the disk. Extents optimize sequential data block layout and reduce the amount of time required for applications to perform sequential read operations, such as database scans.

However, using extents increases write overhead. In the case of randomly writing one data block, when extents are enabled Data ONTAP reads three additional blocks and writes three additional blocks.

Next topics

[When to enable extents](#) on page 331

[Enabling and disabling extents](#) on page 332

When to enable extents

Consider enabling extents when you want to improve the performance of Exchange database validation. However, if increased write overhead is a concern, you might not want to enable extents.

The benefits of enabling extents include the following:

- On volumes that contain only Microsoft Exchange data, enabling extents might improve the performance of Exchange database validation.
- On workloads that perform many small random writes followed by sequential reads, enabling extents might improve sequential read performance.

The costs of enabling extents include the following:

- Enabling extents results in a higher load on the storage system, thereby increasing the latency of other work, especially write latency under a heavy load.
- Unless `vol options vol-name extent` is set to `space_optimized`, enabling extents causes some data in Snapshot copies to be duplicated in the active file system, and it also causes SnapMirror updates to transfer more information, thereby using more space to store the same amount of data.

If you want to enable extents but storage space is a concern, you can enable extents on FlexVol volumes by setting `vol options vol-name extent` to `space_optimized` (instead of `on`). Setting the option to `space_optimized` conserves space but results in degraded read performance through the Snapshot copies. Therefore, if fast read performance through Snapshot copies is a higher priority to you than storage space, do not use `space_optimized`.

Extents might conflict with deduplication by adding new blocks that were previously consolidated during the deduplication process. A deduplication scan might also consolidate blocks that were previously rearranged by extents, thus separating chains of blocks that were sequentially laid out on disk. Therefore, because enabling extents does not predictably optimize sequential data block layout when used on deduplicated volumes, it is best not to enable extents on deduplicated volumes. If you want Microsoft Exchange data to benefit from extents, store it on volumes that are not enabled for deduplication.

The extents options are not supported on FlexCache volumes. If file fragmentation is a concern, enable extents on the original server volume.

Enabling and disabling extents

You can enable or disable extents on a traditional or FlexVol volume.

Step

1. Enter the following command:

```
vol options vol-name extent [on | space_optimized | off]
```

- `on` enables extents for the volume.
Enabling extents might help workloads if you perform many small random writes followed by large sequential reads. However, enabling extents might increase the amount of disk operations performed on the storage system.
- `space_optimized` also enables extents but can be used only on FlexVol volumes.
Using `space_optimized` might be useful if the FlexVol volume has Snapshot copies or is a SnapMirror source. When you use `space_optimized`, the extent update does not result in duplicated Snapshot copies in the active file system, thus conserving space in the volume.

Also, `space_optimized` might reduce the amount of data that SnapMirror needs to move on the next update. However, `space_optimized` might result in degraded Snapshot read performance.

`space_optimized` is not supported if `vol-name` is in an aggregate that was either created prior to Data ONTAP 7.2 or once reverted to a version earlier than Data ONTAP 7.2.

- `off` disables extents for the volume. By default, extents are disabled.

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

Troubleshooting tools

If you experience problems with your storage system, some tools are available to help you understand and avoid problems.

Next topics

[Storage system panics](#) on page 335

[Error messages](#) on page 336

[How to use the NOW site for help with errors](#) on page 337

[How to use the remote management device to troubleshoot the system](#) on page 338

Storage system panics

If your storage system has a serious problem, such as a problem with the hardware or a severe bug in the system software, it might panic.

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.

Reacting to storage system panics

If your storage system panics, there are some steps you can follow to help technical support troubleshoot the problem more quickly.

About this task

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

Steps

1. Access the panic message on the console messages or in the `/etc/messages` file.
2. From the NOW site, navigate to the Panic Message Analyzer tool.
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. Otherwise, call technical support.

Related information

The NOW site - <http://now.netapp.com/>

Error messages

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.

The error message can be logged to the console, a file, or to a remote system, depending on how you have configured message logging.

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.

Next topics

[Using the Syslog Translator to get more information about error messages](#) on page 336

[Accessing the Syslog Translator using FilerView](#) on page 337

Related tasks

[Configuring message logging](#) on page 162

Using the Syslog Translator to get more information about error messages

Error messages are relatively brief to avoid clogging the error logging system. Some messages have more information available through the Syslog Translator.

Steps

1. Go to the NOW site and select **Technical Assistance & Documentation** and then **Syslog Translator**.
2. In the Software field, select **Data ONTAP**.
3. Cut and paste the error message into the Search String field and click **Translate**.

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

Related information

The NOW site - <http://now.netapp.com/>

Accessing the Syslog Translator using FilerView

You can access the Syslog Translator through FilerView.

Steps

1. From FilerView, select **Filer > Syslog Messages**.

The `/etc/messages` file is displayed.

2. Click any message displayed as a hot link to access the Syslog Translator for that message.

Note: If a message is not listed as a hot link, no further information is available from the Syslog Translator for that message.

How to use 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.

Related information

The NOW site - <http://now.netapp.com/>

How to use the remote management device to troubleshoot the system

You can use the remote management device to troubleshoot the system even if you are not physically co-located with the system.

You can use the remote management device to view system console messages, view system events, dump the system core, and issue commands to power-cycle, reset, or reboot the system.

Related concepts

[How to troubleshoot the storage system with the RLM](#) on page 241

[How to troubleshoot the storage system with the BMC](#) on page 268

Related references

[SP commands for troubleshooting the storage system](#) on page 211

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 <code>telnet</code> and <code>rsh</code> connections from this client, to give permission to this client for mounting the <code>/</code> and <code>/home</code> 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 <i>Data ONTAP 7-Mode Storage Management Guide</i> .
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 <code>/etc/passwd</code> 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.
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 storage system prompt is an example of a Command Line Interface.
client	A computer that shares files on a storage system.
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 high-availability 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 <code>/etc/group</code> 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.
high-availability configuration	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 in a high-availability configuration, each system is also referred to as a node.
high-availability configuration interconnect	Cables and adapters with which the two storage systems in a high-availability configuration are connected and over which heartbeat and WAFL log information are transmitted when both systems are running.
high-availability configuration monitor	Software that administers the relationship of storage systems in the high-availability configuration through the <code>cf</code> command.
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 high-availability configuration 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 high-availability configuration.
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 high-availability configuration.
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 <i>Data ONTAP 7-Mode Storage Management Guide</i> .
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 copies enable 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 high-availability configuration; the opposite of <i>giveback</i> .
takeover mode	The method you use to interact with a storage system when it has taken over its partner. The console prompt indicates when the storage system is in takeover mode.
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.
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	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 <i>Data ONTAP 7-Mode Storage Management Guide</i> .
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

- /etc directory 28, 93
- /etc/hosts.equiv file 80
- /etc/log/auditlog file 163, 164
- /etc/messages file 96, 160
- /etc/rc file 165, 167
- /etc/syslog.conf file
 - configuring message logging in 162
 - file format and parameters of 161
- /etc/usermap.cfg file, character coding of 96
- /home file, contents of 92
- /vol/vol0, root volume 89

3DES, for SecureAdmin 51

A

- administration host, logging in to the SP from 195
- administration hosts
 - adding 79, 81
 - defined 79
 - removing 81
 - use of 79
 - where they are specified 80
- administrative level commands 39
- administrator access, managing 115
- administrator accounts
 - changing the password of (passwd) 142
 - reasons for creating 115
- aggregate Snapshot copy management 152
- aggregates
 - aggr copy command 29
 - aggr status command, description of 279
 - aggregate state, displaying (aggr status) 279
 - disk statistics, displaying (aggr status) 279
 - performance improvements for disk-bound aggregates 313
 - root option 100
- alternative boot modes
 - booting 104
- assigning priorities using FlexShare 302
- audit-log file 115, 163
- authentication
 - public key-based 56
 - with SSH 51

- with SSL 61

AutoSupport

- about 175
- configuring 177
- contents of email 186
- defined 175
- events that trigger e-mail 185
- mail host support for 176
- options 177
- options AutoSupport.option (configures AutoSupport) 177
- reboots and 160
- requirements for 176
- technical support and 175
- testing 182
- testing (options autosupport.doit) 182
- troubleshooting 183
- when system reboots 160
- transport protocol 176

B

- banner message for Telnet sessions 69

BMC

- admin mode command syntax 257
- admin mode commands 257
- advanced command syntax 259
- advanced mode commands 259
- AutoSupport messages 265
- command line interface (CLI) 255
- description of 245
- displaying information in admin mode 261
- displaying information in advanced mode 259
- features 247
- firmware update problems, troubleshooting 272
- how to configure 248
- logging in to 253
- managing with Data ONTAP commands 248, 252
- system console redirection feature 260
- troubleshooting communication problems 269
- troubleshooting configuration problems 269
- troubleshooting connection problems 270
- troubleshooting firmware update problems 272
- troubleshooting hardware problems 270
- using AutoSupport options 253
- booting systems from 108, 111

- system event log 264
- boot device
 - booting the storage systems 103
 - recovering from corrupted image 109
- boot options 103
- booting
 - from alternative boot modes 104
 - from firmware prompt 107
 - from maintenance mode 107
- booting the system
 - remotely 108, 111
- browsers, improving security through 61
- Bugs Online 337

C

- capabilities
 - assignment to users 117
 - definition of 116
 - list of supported types 129
 - modifying others' 123
 - types of 129
- cards, expansion, displaying information about 275
- certificate-authority-signed certificates 61
- certificates
 - domain names and 64
 - generating 62
 - installing 63
 - testing 63
 - types of 61
 - used by SSL protocol 61
- change privileges, file ownership 84
- character coding for configuration files 96
- checksums, displaying information 279
- CIFS
 - accessing /etc directory 97
 - accessing /home directory 99
 - administrator accounts in 115
 - editing configuration files using 96
 - client, requirements to manage storage system 80
- client decryption 51
- clients
 - editing configuration file from 96
 - platforms supported by FilerView 75
 - SecureAdmin supported 51
 - CIFS, requirements 80
 - NFS, requirements 80
- commands
 - AutoSupport.option (sets AutoSupport options) 177
 - date (sets system date and time) 156

- halt (halts the storage system) 112
- license 155
- options autosupport.doit (tests AutoSupport) 182
- passwd (changes administrative user password) 142
- passwd (changes storage system system password) 141
- privilege levels 39
- savecore, what it does 159
- stats 289
- timezone (displays and sets system time zone) 158
- useradmin 115
 - administrative level 39
 - advanced level 39
 - options wafl.root_only_chown (sets file ownerships changes) 84
 - privilege level 39
 - reboot (reboots the storage system) 111
 - RSH command list 73
- CompactFlash cards
 - checking the Data ONTAP version of 110
 - description of 24
- configuration
 - display, using sysconfig 275
 - message logging 160
 - of AutoSupport (options AutoSupport.option) 177
- configuration files
 - /etc 92
 - accessing 35
 - backing up 168
 - backing up and cloning 168
 - cloning 169
 - comparing backups 170
 - editing from CIFS client 96
 - editing from NFS client-setup 95
 - hard limits 94
 - restoring 169
 - within /etc directory 93
- configuration prerequisites, SP 191
- configuration, SP 192
- core files 159
- criticaltime (UPS option) 173

D

- data access management 29, 32
- data migration management 29
- Data ONTAP, check version of 110
- data organization management 28
- data protection 29
- data storage management 28
- DataFabric Manager 23

- date, setting storage system time and 156
- decryption, between client and storage system 51
- default directories 92
- default root aggregate 89
- default root volume 89
- device carrier 23
- diagnostic account 146
- directories, default permissions 92
- disks, displaying statistical information for 275
- displaying volume information (sysconfig -v) 275
- domain names, changing storage system 64
- domainusers
 - definition of 116
 - deleting 137
 - granting access to 122
 - listing 133
- DSA key pair, for SSH 56

E

- e0M 45, 46
- encryption
 - with SSH 51
 - with SSL 61
- encryption algorithms supported by SecureAdmin 51
- error message logging, about 336
- Exchange, performance 331
- extents 331

F

- F-Secure, for SecureAdmin 51
- file ownership change privileges 84
- FilerView
 - accessing storage system through 35, 75
 - description 75
 - supported by client platforms 75
 - Help system defined 78
 - interface 78
- files, configuration 92
- filestats command
 - about 281
 - options for 281
- FlexShare
 - about 299
 - buffer cache policy 304, 305
 - default priority, modifying 306
 - default queue 301
 - io_concurrency options 302
 - priorities, assigning 302

- priorities, removing 306
- volume operations and 300
- when to use 300

FTP

- accessing /etc directory 98
- accessing /home directory 99

G

- generating certificates 62
- groups
 - assigning roles to 125
 - assigning users to 120
 - definition of 116
 - deleting 137
 - listing 133
 - naming requirements 117
 - predefined 124
 - reloading from lclgroups.cfg file 126
 - renaming 126
 - Windows special 117
 - setting maximum auxiliary 127

H

- hard limits, configuration files 94
- HMAC, for SecureAdmin 51
- host keys
 - changing the size of 55
 - determining sizes 53
 - setting 53
 - uses of 51
 - using with SSH 53
 - where stored 53
- hosts
 - definition of 79
- HTTP access to log files 100
- HTTPS 49

I

- increasing cache memory 307
- installing certificates 63
- interface, use of FilerView 78

K

- keys
 - private and public 56

- public-based authentication 56
- session 51
- used by SSH protocol 51

Knowledgebase Solutions 337

L

- LCD, on storage system chassis 24
- lclgroups.cfg file, reloading 126
- licenses 154, 155
- log files, accessing using HTTP or HTTPS 100
- LUN restore 151
- LUNs
 - reallocating to improve performance 318

M

- mail host support for AutoSupport 176
- maintenance mode
 - booting from 107
- man-in-the-middle warning 55
- Manage ONTAP Developer SDK software 35
- message files, accessing using HTTP or HTTPS 100
- message logging, configuring 160
- Microsoft Exchange, performance 331
- mount privileges, controlling of (options
 - nfs.mount_rootonly) 84
- multiprotocol file and block sharing 28

N

- naming requirements for useradmin command 117
- NDMP 29
- Network file service 27
- NFS
 - access to /etc directory 97
 - access to /home directory 98
- NFS client
 - requirements to manage storage system 80
- nonlocal users, granting access to 122
- NVFAIL 29
- NVRAM
 - halt command to save data to disk 112
 - description of 24

O

- obsolete domain names, and SSL 64
- online command-line help 38
- OpenSSH
 - for SecureAdmin 51
 - generating key pairs in 56

- options
 - security 85
- ownership change privileges, file 84

P

- PAM (Performance Acceleration Module) 307
- panics 335
- password rules, changing 143
- passwords
 - changing (passwd) 141
 - managing security using 140
- perfmom, using to monitor performance 298
- performance
 - Microsoft Exchange read 331
 - monitoring with perfmom 298
 - read 330
 - read reallocation 330
- Performance Acceleration Module 307
- performance improvements, in storage systems
 - backup rate 316
 - caching client data to reduce traffic 315
 - disk-bound volume 313
 - large transfer 315
 - maintain adequate free blocks and inodes 314
 - reallocate command 315
 - using TCP 312
 - WAFL extended cache 307
 - balancing NFS traffic on interfaces 312
- permissions of default directories (/etc, /home) 92
- plexes, displaying information about 279
- priorities
 - assigning using FlexShare 302
 - removing using FlexShare 306
- privilege levels for Data ONTAP commands 39
- privileges, file ownership change 84
- public-key encryption 51
- PuTTY, for SecureAdmin 51

Q

- quota file, character coding for 96

R

- RAID
 - displaying statistics (aggr status) 279
 - displaying statistics (sysconfig -r) 275
 - displaying statistics (vol status) 280
- reading files 170

- reallocate commands
 - reallocate off 329
 - reallocate on 321
 - reallocate quiesce 327
 - reallocate restart 327
 - reallocate schedule 323
 - reallocate start 321, 324
 - reallocate start -A 322, 324
 - reallocate status 328
 - reallocate stop 328
 - reallocate schedule -d 324
- reallocation
 - best practices 329
 - defining scans
 - aggregates 322
 - LUNs, files, or volumes 321
 - deleting a scan 328
 - deleting scan schedule 324
 - disabling scans 329
 - enabling scans 321
 - full 325
 - managing scans 320
 - measure-only 326
 - quiescing scans 327
 - restarting scans 327
 - scans 317
 - scheduling scans 323
 - starting one-time scan 324
 - viewing scan status 328
 - with LUNs, files, or volumes 318
 - read 330
- rebooting the system
 - from the console 111
- reinitialization
 - of SSH 55
 - of SSL 64
- Remote LAN Module (RLM) 216
- remote management
 - BMC 247
 - RLM 217
 - SP 189
- Remote Management Controller (RMC) 24
- remote management device 46, 108, 109, 111
- Remote Shell (RSH) 70
- Remote Support Agent (RSA) 273
- remote system management 189
- removing priorities using FlexShare 306
- requirements
 - to manage storage system on NFS clients 80
 - for Web browser 76
- RLM
 - admin mode command syntax 227
 - admin mode commands 227
 - advanced command syntax 229
 - advanced mode commands 229
 - AutoSupport messages 238
 - command line interface (CLI) 231
 - description of 216
 - displaying information in advanced mode 229
 - features 217
 - how to configure 218
 - logging in to 222
 - managing with Data ONTAP commands 218, 231
 - system console redirection feature 229
 - troubleshooting communication problems 242
 - troubleshooting configuration problems 242
 - troubleshooting connection problems 242
 - using AutoSupport options 232
 - booting systems from 108, 111
 - displaying information in admin mode 234
 - down filer events 232
 - down system events 232
 - SNMP traps 232, 233
 - system event log 236
 - troubleshooting hardware problems 243
- RLM (Remote LAN Module) 216
- RMC (Remote Management Controller) 24
- roles
 - assigning to groups 125
 - creating 131
 - definition of 116
 - deleting 137
 - listing 133
 - modifying 132
 - naming requirements 117
 - predefined 128
- root option for aggregates 100
- root password, changing 141
- root volume
 - changing 100
 - default name 89
 - directories contained within 92
 - space guarantees and 91
 - minimum size 91
 - size requirement 91
- RSA key pair
 - definition of 56
 - generating for SSH 1.x 56
 - generating for SSH 2.0 57
 - where stored 56, 57

RSA/DSA, for SecureAdmin 51

RSH (Remote Shell)

access to storage system 70

using with Windows 73

RSH commands

accessing storage system from a PC client 73

accessing storage system from a UNIX client 72

displaying session information 74

list of 73

privilege levels 39

use with user names and passwords 71

S

scans, reallocation 320

secure connection, testing 63

Secure FilerView, improving security using 61

secure protocols 49

secure session, creating with SecureAdmin 51

Secure Shell (SSH) 49

Secure Sockets Layer (SSL) 49

Secure Sockets Layer (SSL) protocol

SSLv2 65

SSLv3 65

SecureAdmin

authentication supported 51

creating a secure session with 51

displaying status of 66

encryption algorithms supported 51

improving security with SSH 51

improving security with SSL 61

managing SSH portion 53

managing SSL portion 61

security

improving using Secure FilerView 61

improving using SecureAdmin 49

improving using SSH 51

limiting Telnet access 82

passwords, managing 140

controlling file ownership changes (options

waf.root_only_chown) 84

controlling mount privileges (options

nfs.mount_rootonly) 84

password options 143

settings 50

self-signed certificates 61

server keys

changing the size of 55

setting 53

size guidelines for 53

uses of 51

using with SSH 53

Service Processor (SP) 189

session keys, uses of 51

slots, expansion (storage system hardware) 25

SnapMirror 29

SnapRestore 29

Snapshot copy, aggregate 152

SnapVault 29

SP

admin mode commands 199

advanced mode commands 202

AutoSupport messages 213

command line interface (CLI) 197, 203

description of 189

how to configure 191, 192

logging in to 194

logging in to from an administration host 195

managing with Data ONTAP commands 191, 203

system console redirection feature 197, 202

system event log 212

troubleshooting connection problems 215

updating firmware 214

using AutoSupport options 204

accessing from system console 196

booting systems from 108, 111

commands for troubleshooting 211

console log 213

down system events 205

managing with Data ONTAP commands 191, 203

online help 198

prerequisites for configuration 191

sensors 206

sensors, discrete 208–210

sensors, threshold-based 206

SNMP traps 205, 206

SP (Service Processor) 189

special system files

.bplusvtoc_internal 151

.vtoc_internal 151

SSH (Secure Shell) commands

secureadmin disable ssh 55, 56

secureadmin enable ssh 55, 56

secureadmin setup -f ssh 55

secureadmin setup ssh 53

secureadmin status 66

SSH (Secure Shell) protocol

authentication with 51

creating a secure session with 51

determining host and server key size using 53

- disabling or enabling 56
 - encryption with 51
 - host keys 53
 - improving security with 51
 - keys used by 51
 - managing 53
 - reinitializing 55
 - server keys 53
 - setting up and starting 53
- SSH Communications Security client, for SecureAdmin 51
- SSH interactive
 - configuring a timeout period 70
 - controlling the timeout period 69
- SSL (Secure Sockets Layer) commands
 - secureadmin addcert ssl 63
 - secureadmin disable all 66
 - secureadmin disable ssl 64
 - secureadmin enable all 66
 - secureadmin enable ssl 64
 - secureadmin setup ssl 62
 - secureadmin status 66
- SSL (Secure Sockets Layer) protocol
 - authentication with 61
 - certificates used with 61
 - enabling or disabling 64
 - improving security with 61
 - managing 61
 - reinitializing 64
 - setting up and starting 62
- startup commands 165
- statistics commands
 - aggr status command, description of 279
 - checking expansion cards 275
 - displaying aggregate state statistics 279
 - displaying chassis environment status 286
 - displaying Data ONTAP version 275
 - displaying disk information
 - aggr status 279
 - vol status 280
 - displaying Fibre Channel driver statistics 288
 - displaying link statistics 288
 - displaying overall storage system information 275
 - displaying RAID and checksum information 275, 279, 280
 - displaying relative environment information 286
 - displaying relative physical drive position 288
 - displaying shelf environment status 286
 - displaying tape drive information 275
 - displaying tape library information 275
 - displaying volume
 - language (vol status) 280
 - displaying volume state statistics 280
 - environment description of 285
 - Fibre Channel statistics, description of 288
 - SAS statistics, description of 288
 - sysconfig command
 - description of 275
 - vol status command, description of 280
- stats command
 - about 289
 - background mode 294
 - controlling output 295
 - counters 289
 - instances 289
 - objects 289
 - preset files 297
 - repeat mode 293
 - singleton mode 292
- status commands
 - aggr status -d (displays disk statistics) 279
 - aggr status -r (displays RAID statistics) 279
 - aggr status (displays aggregate state) 279
 - environment chassis (displays shelf environment information) 286
 - environment command, description of 285
 - environment status (displays all storage system environment information) 286
 - environment status shelf (displays shelf environment information) 286
 - fcstat device_map (displays relative physical drive position) 288
 - fcstat fcal_stats (displays fibre channel driver statistics) 288
 - fcstat link_stats (displays link statistics) 288
 - filestats command, description of 281
 - sasadmin (displays SAS adapter and expander information) 288
 - sasstat adapter_state (displays state of a logical adapter) 288
 - sasstat dev_stats (displays statistics for disk drives connected to SAS channels) 288
 - sasstat expander (displays SAS expander configuration) 288
 - sasstat expander_map (displays SAS expander product information) 288
 - sasstat expander_phy_state (displays SAS expander physical state) 288
 - sasstat shelf (displays pictorial representation of the drive population of a shelf) 288

U

Uninterruptible Power Supply (UPS) 172

UPS

- adding a device to be monitored 287
- enabling or disabling monitoring of 287
- management 172
- shutdown process 173

user account, changing password for 142

useradmin

- examples 138
- naming requirements 117

users

- assigning to groups 120
- changing passwords 142
- creation examples 138
- definition of 116
- deleting 137
- examples of creating 138
- listing 133
- managing root access 118
- modifying capabilities of 123
- naming requirement 117

V

Vandyke SecureCRT, for SecureAdmin 51

version checking, Data ONTAP 110

volumes

- disk statistics, displaying (vol status) 280
- vol status command, description of 280
- volume language, displaying (vol status) 280
- volume state, displaying (vol status) 280
- volume statistics, displaying 275
- vol copy 29

W

WAFL (Write Anywhere File Layout) 23, 307

WAFL extended cache

- about 307
- buffer cache policies 310
- displaying configuration 311
- displaying usage and access information 311
- low-priority user data blocks 308
- normal user data blocks 308
- system metadata cache 309

warnings

- man-in-the-middle 55
- obsolete domain names 64

warningtime (UPS option) 173

Web browser requirements 76

Windows

- administrator accounts in 115
- domainusers, granting access to 122
- network commands 35
- special groups 117

Write Anywhere File Layout (WAFL) 23, 307

writing files 170

