




Data ONTAP® 8.2

System Administration Guide

For 7-Mode



NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089
U.S.

Telephone: +1(408) 822-6000
Fax: +1(408) 822-4501
Support telephone: +1(888) 4-NETAPP
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 210-06052_A0
April 2013

Contents

Introduction to NetApp storage	12
Main components of a storage system	12
Internal components of a storage system	12
Slots and serial ports of a storage system	13
What disk shelves do	14
Third-party storage	14
Components of a Data ONTAP-v storage system	15
Key features for Data ONTAP	15
Network file services	15
Multiprotocol file and block sharing	16
Data storage management	16
Data organization management	17
Data access management	17
Data migration management	17
Data protection	18
System management	20
AutoSupport	21
How to interface with Data ONTAP	22
Methods for administering a storage system	22
Methods for administering a Data ONTAP-v storage system	23
Data ONTAP command-line interface	23
Displaying command history	24
Using the command-line editor	24
How to use online command-line help	25
Data ONTAP commands at different privilege levels	25
How different privilege settings apply to different sessions	26
Initial privilege level	26
Setting the privilege level	26
How to access the storage system	28
Methods for accessing a storage system	28
Methods for administering the system	28
Controlling the sharing of a console session	28

Rules that apply to console, Telnet, and SSH-interactive sessions	29
What the e0M interface is	31
How to access a storage system from the console	33
Accessing the storage system by using the serial port	33
Accessing the system by using the remote management device	34
Secure protocols and storage system access	35
The default security settings	35
Understanding the SSH protocol	36
The SSL protocol	46
Determining whether secure protocols are enabled	52
Enabling or disabling secure protocols	52
How to access a storage system by using Telnet	52
Starting a Telnet session	53
Terminating a Telnet session	54
Configuration for Telnet sessions	54
How to access a storage system by using a Remote Shell connection	56
When to use RSH commands with user names and passwords	56
Accessing a storage system from a UNIX client by using RSH	57
Accessing a storage system from a Windows client by using RSH	58
Commands not accepted when using RSH	59
How to reset options to default values from RSH	60
Displaying RSH session information	60
Understanding OnCommand System Manager	61
How to manage access from administration hosts	62
Reasons to designate a workstation as an administration host	62
Administration host privileges	63
Requirements for using a client	63
How to specify administration hosts	63
Adding administration hosts	64
Removing administration hosts	64
Methods for controlling storage system access	65
Controlling Telnet access using host names	65
Restricting protocol access	65
Controlling the NFS mount privilege	66
Controlling file ownership change privileges	67
Controlling anonymous CIFS share lookups	67

Options that help maintain security	68
Allowing only secure access to the storage system	69
Managing the root aggregate and the root volume	71
Understanding the root aggregate	71
Understanding the root volume	71
Recommendations for the root volume	72
Sizing considerations for root FlexVol volumes	73
Default directories in the root volume	74
Permissions for the default directories	74
The /etc directory	75
How to access the default directories on the storage system	79
Accessing the /etc directory from an NFS client	79
Accessing the /etc directory from a CIFS client	79
Accessing the /etc directory with FTP	79
Accessing the /etc directory with SFTP	80
Accessing the /home directory from an NFS client	81
Accessing the /home directory from a CIFS client	81
Accessing the /home directory with FTP	81
Accessing the /home directory with SFTP	82
Accessing log files using HTTP or HTTPS	82
Changing the root volume	83
Starting and stopping the storage system	86
How to boot the storage system	86
Ways to boot the storage system	86
Rebooting the storage system at the system prompt	87
Booting Data ONTAP at the boot environment prompt	87
Rebooting the storage system remotely	88
Recovering from a corrupted image of the system's boot device	88
Checking available Data ONTAP versions	89
About rebooting the storage system	90
Rebooting the storage system from the system console	90
Rebooting the storage system remotely	90
Halting the storage system	91
Managing the storage system by using the boot menu	92
How to manage administrator and diagnostic access	95
Reasons for creating administrator accounts	95

What users, groups, roles, and capabilities are	95
How users are assigned capabilities	96
Requirements for naming users, groups, and roles	96
Windows special groups	97
About changing capabilities of other groups and roles	97
Root access to the storage system	97
Disabling root account access to the storage system	98
Displaying the status of root access	98
How to manage users	99
Creating users and assigning them to groups	99
Granting access to Windows domain users	101
How to grant permissions for MMC	102
About changing another user's capabilities	102
How to manage groups	103
Predefined groups	103
Assigning roles to groups by creating or modifying a group	104
Renaming a group	105
Loading groups from the lclgroups.cfg file	105
Setting the maximum number of auxiliary UNIX groups allowed for a user	106
How to manage roles	106
Predefined roles	106
Supported capability types	108
Creating a new role and assigning capabilities to roles	110
Modifying an existing role or its capabilities	111
Users, groups, and roles	112
Commands that list users, domain users, groups, or roles	112
Commands that delete users, domain users, groups, or roles	117
Administrative user creation examples	118
Example of creating a user with custom capabilities	118
Example of creating a user with no administrative capabilities	119
Granting users in LDAP groups access to the system and mapping them to specified roles	119
How to manage passwords for security	121
Changing the storage system password	123
Changing a local user account password	124

Data ONTAP options for managing password rules	124
Uses of the systemshell and the diagnostic account	128
Enabling or disabling the diagnostic account	128
Setting the password for the diagnostic account	129
Accessing the systemshell	130
General system maintenance	132
Special system files	132
Managing aggregate Snapshot copies	132
Considerations for increasing the aggregate Snapshot reserve	133
Managing licenses	134
License types	134
Commands for managing licenses	135
Setting the system date and time	136
Synchronizing the system time	137
Data ONTAP options for managing system time	137
Displaying and setting the system time zone	138
Managing core dump files	139
Considerations for managing core dump files	139
Methods of segmenting core dump files	139
Commands for managing core segmenting	140
Automatic technical support notification upon system reboots	141
Understanding message logging	142
The /etc/syslog.conf file	142
Sample /etc/syslog.conf file	144
Configuring message logging	144
Understanding audit logging	145
Configuring audit logging	146
Enabling or disabling read-only API auditing	146
Startup configuration for the storage system	146
Commands in the /etc/rc file	146
Editing the /etc/rc file	148
Recovering from /etc/rc errors	149
Storage system configuration backup and cloning	149
Backing up a storage system configuration	150
Cloning a storage system configuration	151
Restoring a storage system configuration	151

Comparing storage system configurations and backup configuration files	152
About writing and reading files on the storage system	153
Writing a WAFL file	153
Reading a WAFL file	154
Monitoring the storage system	155
Managing event messages	155
Displaying event information	155
Displaying event log information	156
Managing AutoSupport	157
When and where AutoSupport messages are sent	157
How event-triggered AutoSupport messages work	159
How AutoSupport On Demand obtains delivery instructions from technical support	160
What data AutoSupport messages contain	161
Structure of AutoSupport messages sent via email	163
AutoSupport severity types	164
AutoSupport transport protocols	164
Setting up AutoSupport	166
Getting AutoSupport message descriptions	168
Commands for managing AutoSupport	168
AutoSupport options	169
What the AutoSupport manifest is	176
Troubleshooting AutoSupport	176
Monitoring the health of your system	181
How health monitoring works	181
What health monitors are available	183
Responding to degraded system health	183
Commands for monitoring the health of your system	186
Managing a storage system remotely	189
Managing a system remotely by using the Service Processor	189
Configuring the SP network	191
Accounts that can access the SP	193
Accessing the SP from an administration host	194
Accessing the SP from the serial console	195
Accessing the serial console from the SP	196

SP CLI and system console sessions	197
Using online help at the SP CLI	197
Commands for managing a system at the SP admin privilege level	198
Commands for managing a system at the SP advanced privilege level	201
How to determine the status of a threshold-based SP sensor	202
How to determine the status of a discrete SP sensor	204
Troubleshooting a system by using the SP	206
Managing the SP with Data ONTAP	207
Managing a system remotely by using the Remote LAN Module	215
What the RLM does	216
Ways to configure the RLM	217
Prerequisites for configuring the RLM	218
Configuring the RLM	218
Accounts that can access the RLM	221
Restricting RLM access to only the specified administration hosts	223
Configuring automatic logout of idle SSH connections to the RLM	224
Logging in to the RLM from an administration host	225
Accessing the serial console from the RLM	226
RLM CLI and system console sessions	227
Using online help at the RLM CLI	227
Commands for managing the RLM at the admin privilege level	228
Commands for managing the RLM at the advanced privilege level	230
Troubleshooting the storage system by using the RLM	231
Managing the RLM with Data ONTAP	232
Troubleshooting RLM connection problems	235
System information	236
Displaying storage system configuration information	236
Displaying aggregate information	238
Displaying volume information	239
Commands for displaying environmental status	240
Getting Fibre Channel information	241
Getting SAS adapter and expander information	241
Storage system information and the stats command	242
Viewing the list of available counters	243
Getting detailed information about a counter	244
Using the stats command interactively in singleton mode	245

Using the stats command interactively in repeat mode	246
Collecting system information with the stats command in background mode	247
Changing the output of a stats command	248
About the stats preset files	250
How to get system information using perfmon	251
How to get system information using perfstat	251
Managing system performance	252
Managing storage system resources by using FlexShare	252
What FlexShare is	252
When to use FlexShare	253
FlexShare and priority levels	254
Considerations for using FlexShare in storage systems with a high- availability configuration	254
How the default FlexShare queue works	255
FlexShare and the global io_concurrency option	255
FlexShare and the buffer cache policy values	255
Using FlexShare	256
Increasing WAFL cache memory	260
How Flash Pools and Flash Cache compare	260
Enabling and disabling WAFL external cache	261
Caching normal user data blocks	262
Caching low-priority user data blocks	262
Caching only system metadata	263
Displaying the WAFL external cache configuration	263
Displaying usage and access information for WAFL external cache	264
Preserving the cache in the Flash Cache family of modules	264
Optimizing LUN, file, volume, and aggregate layout	267
What a reallocation scan is	267
Reasons to use LUN, file, or volume reallocation scans	268
Reasons to use aggregate reallocation scans	268
Reasons to use physical reallocation scans	269
How a reallocation scan works	269
Managing reallocation scans	270
How to use reallocation scans most efficiently	280
Improving read performance	281

What read reallocation is	281
Enabling or disabling read reallocation	281
Improving write performance	282
How free space reallocation optimizes free space	282
When to enable free space reallocation	283
When to use free space reallocation with other reallocation features	284
How free space reallocation differs from an aggregate reallocation scan ...	284
Types of aggregates that free space reallocation can and cannot optimize .	284
Commands for managing free space reallocation	284
Troubleshooting tools	286
Storage system panics	286
Reacting to storage system panics	286
Error messages	287
Using the Syslog Translator to get information about error messages	287
How to use the NetApp Support Site for help with errors	288
How to use the remote management device to troubleshoot the system	288
Glossary	289
Copyright information	297
Trademark information	298
How to send your comments	299
Index	300

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 disk arrays. They provide a modular hardware architecture running the Data ONTAP operating system and WAFL (Write Anywhere File Layout) software.

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 System Manager, and through remote management devices such as the Service Processor (SP) and the Remote LAN Module (RLM).

For information about all of the models of NetApp storage systems, see the NetApp Products page.

Related information

NetApp Support Site: support.netapp.com

Main components of a storage system

A storage system running Data ONTAP has a main unit, which is the 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 consists of the following main components:

- The storage controller, which is the component of a storage system that runs the Data ONTAP operating system and controls its disk subsystem
- The disk shelves, which contain disk drives and are attached to a storage system
V-Series systems fulfill client requests from either disk shelves or logical units on the back-end storage arrays from IBM, Hitachi, HP, EMC, and more. See the *Interoperability Matrix* for information about the storage arrays that V-Series supports.

Related information

NetApp Support Site: support.netapp.com

Internal components of a storage system

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

The following table describes the internal components of a storage system:

Component	Description
Controller	The controller is the component that runs the Data ONTAP operating system and controls its disk subsystem.
System memory	System memory stores information temporarily.
Nonvolatile RAM (NVRAM) or nonvolatile memory (NVMEM)	Data ONTAP uses nonvolatile memory (NVRAM or NVMEM, depending on the platform) to log network transactions as a data integrity measure. In case of a system or power failure, Data ONTAP uses the contents of the nonvolatile memory 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.
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 device, such as the Service Processor (SP) or the Remote LAN Module (RLM)	The remote management device provides remote platform management capabilities for the storage system. It enables remote access to the storage system console over a network regardless of the operating state of the storage system. It also monitors and maintains hardware event logs for the storage system and generates alerts based on system status.

Slots and serial ports of a storage system

The storage system has slots for external connections and serial ports for a console and diagnostic hardware.

The following table describes the slots and serial 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 • Flash Cache family of modules and Performance Acceleration Modules (PAM) • Nonvolatile memory adapters

Component	Description
serial ports	<p>The serial ports include the following:</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

For information about how to configure host adapters for your storage system, see the *Hardware Universe* (formerly the *System Configuration Guide*) at support.netapp.com/knowledge/docs/hardware/NetApp/syscfg/index.shtml.

What disk shelves do

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 Storage Management Guide for 7-Mode*.

For information about disk shelves connected to V-Series systems, see 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 *Interoperability Matrix*.

For information about setting up a specific storage array to work with Data ONTAP, see the *V-Series Implementation Guide for Third-Party Storage*.

Components of a Data ONTAP-v storage system

A storage system based on Data ONTAP-v technology is a software-based data storage and retrieval system.

The software version of a NetApp storage controller executes Data ONTAP within a virtual machine on a host server. This enables Data ONTAP to run on servers supported by VMware vSphere to manage storage in virtualized environments.

The virtual storage system includes:

- A storage system main unit and storage disks collocated within a single chassis.
- Data ONTAP software stored on the boot device located on a storage disk.
- Network access to Data ONTAP without the use of dedicated ports.
- The Data ONTAP CLI and System Manager to perform storage management.
- The Data ONTAP-v Administration tool (dvsadmin) to perform system management.

The virtual storage system does not include dedicated remote management devices. Management is done using the network.

Monitoring, troubleshooting, logging, reporting, and collection of environmental and other system information are performed by the host server.

Key features for Data ONTAP

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.

Network file services

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 their clients. Supported protocols for NAS are NFS, CIFS, HTTP, and FTP for file access. Supported protocols for SAN (block) are iSCSI, FC, and FCoE for block-storage access.

Client workstations are connected to the storage system through direct-attached or TCP/IP network-attached connections, or through FC, fabric-attached connections.

For information about configuring a storage system in a NAS network, see the *Data ONTAP Network Management Guide for 7-Mode*.

For information about configuring a storage system in a SAN fabric, see the *Data ONTAP SAN Administration Guide for 7-Mode*.

Multiprotocol file and block sharing

You can use several protocols to access data on the storage system.

You can use the following protocols for sharing files or blocks:

- 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
- FC (Fibre Channel)—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 File Access and Protocols Management Guide for 7-Mode*.

For information about SAN block access protocols, see the *Data ONTAP SAN Administration Guide for 7-Mode*.

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.

For more information about data storage management, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Data organization management

Data ONTAP organizes the data in user and system files and volumes, optionally in qtrees, and, for SAN environments, in Logical Unit Numbers (LUNs). Aggregates provide storage to the volumes that they contain.

For more information about data organization management, see the *Data ONTAP Storage Management Guide for 7-Mode* and the *Data ONTAP SAN Administration Guide for 7-Mode*.

Related concepts

[Understanding the root volume](#) on page 71

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 File Access and Protocols Management Guide for 7-Mode*.
- Creates 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, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

Data migration management

Data ONTAP enables you to manage data migration in several ways.

You can use the following Data ONTAP functionality to manage data migration:

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

Data protection

Data ONTAP provides a wide range of data protection capabilities, such as `aggr copy`, MetroCluster, NDMP, NVFAIL, SnapLock, SnapMirror, SnapRestore, Snapshot copies, SnapVault, SyncMirror, tape backup and restore, virus scan support, and `vol copy`.

Feature	Description
<code>aggr copy</code>	<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 Storage Management Guide for 7-Mode</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 High Availability and MetroCluster Configuration Guide for 7-Mode</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 Data Protection Online Backup and Recovery Guide for 7-Mode</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 Data Protection Online Backup and Recovery Guide for 7-Mode</i>.</p>
SnapLock software	<p>SnapLock provides an alternative to traditional optical WORM (write-once-read-many) storage systems for nonrewritable data.</p> <p>For information about SnapLock, see the <i>Data ONTAP Archive and Compliance Management Guide for 7-Mode</i>.</p>
SnapMirror software	<p>System-to-system Snapshot mirroring enables you to mirror Snapshot copies on one storage system to a partner system. If the original storage system is 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 Data Protection Online Backup and Recovery Guide for 7-Mode</i>.</p>

Feature	Description
SnapRestore software	<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 Data Protection Online Backup and Recovery Guide for 7-Mode</i>.</p>
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 Storage Management Guide for 7-Mode</i>.</p> <p>For information about Snapshot copies, see the <i>Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode</i>.</p>
SnapVault software	<p>SnapVault combines Snapshot schedules and Qtree SnapMirror to provide disk-based data protection for storage systems.</p> <p>Using SnapVault, you can periodically replicate selected Snapshot copies from multiple client 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 Data Protection Online Backup and Recovery Guide for 7-Mode</i>.</p>
Storage Encryption	<p>Storage Encryption protects your data at rest by storing it encrypted on the disk. In a standard storage environment, data is written to disk in cleartext format. This makes the data vulnerable to potential exposure to unauthorized users when disks removed from a storage system are lost or stolen. Storage Encryption is an optional feature that you can enable for additional data protection. It is available on certain supported storage controllers and disk shelves that contain disks with encryption functionality. It does not require a separate license key. It happens without a perceptible disk performance decrease or boot time increase. The only additional hardware required is an external key management server. The self-encrypting disks in the storage system automatically encrypt the data for storage. When you enable Storage Encryption, the disks require authentication to access and decrypt the data. The authentication key is stored securely on an external key management server that is linked to the storage system.</p>

Feature	Description
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 controller. 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 Storage Management Guide for 7-Mode</i>. For information about SyncMirror, see the <i>Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode</i>.</p>
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 Data Protection Tape Backup and Recovery Guide for 7-Mode</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 Data Protection Online Backup and Recovery Guide for 7-Mode</i>.</p>
<code>vol copy</code>	<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 Data Protection Online Backup and Recovery Guide for 7-Mode</i>.</p>

System management

Data ONTAP enables you to manage system activities and resources.

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 MultiStore Management Guide for 7-Mode*.

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

AutoSupport

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

Related concepts

[*Managing AutoSupport*](#) on page 157

How to interface with Data ONTAP

You interface with Data ONTAP to administer your storage system.

Methods for administering a storage system

You can use Data ONTAP, the remote management device, Windows, configuration files, OnCommand System Manager, or the NetApp Manageability 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 enables 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 in the following ways:

 - 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 or 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 the Windows operating system

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 by 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, and then editing the file from the client.

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

OnCommand System Manager is a web-based graphical management interface that enables you to manage storage systems and storage objects, such as disks, volumes, and aggregates. For more information about OnCommand System Manager, see the NetApp Support Site.

- **NetApp Manageability SDK**

The NetApp Manageability SDK software contains resources necessary to develop third-party applications that monitor and manage storage systems. The NetApp Manageability SDK is available to all users of the NetApp Support Site 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 that runs on Linux or Solaris and simulates the NetApp storage system to a very low level is also available as a separate distribution. For more information, see the NetApp Manageability SDK page.

Related concepts

Managing a storage system remotely on page 189

Default directories in the root volume on page 74

Understanding OnCommand System Manager on page 61

Related information

NetApp Support Site: support.netapp.com

NetApp Manageability SDK: communities.netapp.com/docs/DOC-1152

Methods for administering a Data ONTAP-v storage system

You interface with the Data ONTAP-v administration tool (dvadmin) in order to install, configure, and manage the Data ONTAP-v storage software.

Because there is no physical serial port on the Data ONTAP-v virtual machine, you connect to both Data ONTAP-v and Data ONTAP over the network.

There are multiple ways to configure and manage a Data ONTAP-v storage system:

- Connect to the Data ONTAP-v virtual machine over the network and use the dvadmin CLI to manage the virtual machine configuration.
- Connect to Data ONTAP over the network and use the storage system CLI or System Manager to manage the storage system.

See the *Data ONTAP Edge Installation and Administration Guide* for more information.

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 enclose that element in quotation marks, as shown in the following 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 ignores any information following the #.

Displaying command history

Data ONTAP 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.

Step

1. Use the applicable key combination to move the cursor within the same line and edit the command:

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

If you want to...	Then press...
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 that you must choose one of the 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 25

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.

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, or a secure shell client application.

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.

Methods for administering the system

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

You can use the following methods:

- Through a console that is attached by a cable to the storage system's serial port
- Through the Ethernet network interface card (NIC) that is preinstalled in the storage system
The NIC enables you to connect to a TCP/IP network and administer the storage system from a client that uses the following:
 - A Telnet session
 - A Remote Shell connection
 - A web browser
 - 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 sessions, the command fails if SSH is configured to automatically send an End-Of-File (EOF) or used with the option `-n`.

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 68

[Predefined groups](#) on page 103

[Predefined roles](#) on page 106

[Supported capability types](#) on page 108

Related tasks

[Creating users and assigning them to groups](#) on page 99

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.

What the e0M interface is

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) is connected to an internal Ethernet switch. The internal Ethernet switch provides connectivity to the e0M interface and the remote management device, such as the SP, or the RLM. 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 Software Setup Guide for 7-Mode*.

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

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

The e0M is a low bandwidth interface that should be configured for management traffic only. Running data traffic over the e0M interface can cause performance degradation and routing problems. For information about blocking data transfer over the e0M interface, see the *Data ONTAP Network Management Guide for 7-Mode*.

Related concepts

[*Managing a system remotely by using the Remote LAN Module*](#) on page 215

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 session on a client by using a management protocol such as SSH.
For information about how to use the e0M interface with SNMP, see the *Data ONTAP Network Management Guide for 7-Mode*.
2. Connect to the storage system by using the IP 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.

Related concepts

[*How to access a storage system by using Telnet*](#) on page 52

[*How to access a storage system by using a Remote Shell connection*](#) on page 56

[*How to manage SSH*](#) on page 38

[*The default security settings*](#) on page 35

Differences between the e0M interface and the remote management device

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.) However, the e0M interface serves as the dedicated interface for management traffic, and 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 (which can be the SP or the RLM, depending on the storage system model) 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.

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.

Related concepts

[Rules that apply to console, Telnet, and SSH-interactive sessions](#) on page 29

Accessing the storage system by using the serial port

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

About this task

If the values of the following options are changed, the new values take effect only after the console session is reestablished:

- `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.

Steps

1. At the console, press Enter.
The system responds with the login prompt.
2. At the login prompt, do one of the following:

To access the storage system with...	Enter the following account name...
The default system account	<code>root</code>
An alternative administrative user account	<code>username</code>

- The system responds with the password prompt.
3. Enter the password for the root or administrative user account, and then press Enter.
 4. If 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.

Accessing the system by using the remote management device

You can access the system remotely by using the remote management device. Depending on your storage system model, the remote management device can be the SP or the RLM.

Steps

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

```
ssh remote_management_username@IP_for_remote_management_device
```

remote_management_username is the “naroot” account or a Data ONTAP user account with the credentials of the “admin” role or a role with the `login-sp` capability.

The 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. When the system displays the login prompt, enter an appropriate account name for the storage system:

If you are using...	Enter the following account name...
The system root account	root
An administrative user account	storage_system_username
Note: storage_system_username is an administrative user account for the storage system.	

- 4. Enter the password for the account, and then press Enter.

The storage system prompt appears.

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

Related concepts

Managing a storage system remotely on page 189

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 Data ONTAP APIs.

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 69

Understanding the SSH protocol

The Secure Shell (SSH) protocol performs public-key encryption using a host key and a server key. SSH improves security by providing a means for the storage system to authenticate the client and by generating a session key that encrypts data sent between the client and storage system.

The SSH server version running on Data ONTAP is Data ONTAP SSH version 1.0. For information about the Common Vulnerabilities and Exposures (CVE) fixes implemented in Data ONTAP, see the Suspected Security Vulnerabilities page on the NetApp Support Site.

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

Data ONTAP supports the following SSH clients:

- OpenSSH client version 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.	

Note: Some characters, such as question mark (?), period (.), asterisk (*), and caret (^), 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, you can use an escape sequence before the character (`ssh ip_address \?`) or enclose the character in quotes (`ssh ip_address ' ? '`).

Data ONTAP supports password authentication and public-key-based authentication. It 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

Related information

[Suspected Security Vulnerabilities page: support.netapp.com/NOW/knowledge/docs/olio/scanner_results](http://support.netapp.com/NOW/knowledge/docs/olio/scanner_results)

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 95

[The default security settings](#) on page 35

Related tasks

[Restricting protocol access](#) on page 65

Setting up and starting SSH

The SSH setup process involves creating host and server keys.

About this task

You can determine the size of the host and server keys by using the following guidelines:

- If you are using SSH 1.x, the size of the host and server keys can range from 384 bits to 2,048 bits.
- If you are using SSH 2.0, 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 SSH 1.x, the host key is stored in the `/etc/sshd/ssh_host_key` file.

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

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 SSH 2.0, 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 SSH 1.x host key.

The default size for the host key is 768 bits.

3. When prompted, enter a size for the SSH 1.x server key.

The default size for the server key is 512 bits.

4. When prompted, enter a size for the SSH 2.0 host keys.

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.

Result

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

```
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
@          WARNING: HOST IDENTIFICATION HAS CHANGED!          @
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
IT IS POSSIBLE THAT SOMEONE IS DOING SOMETHING NASTY!
Someone could be eavesdropping on you right now (man-in-the-middle attack)!
It is also possible that the host key has just been changed.
```



```
Please contact your system administrator.
Add correct host key in /u/sisa/.ssh/known_hosts to get rid of this message.
Agent forwarding is disabled to avoid attacks by corrupted servers.
Are you sure you want to continue connecting (yes/no)?
```

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 39

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.

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/sshhd/user_name/.ssh/authorized_keys` file.

Examples of generating an RSA key pair

The following is an example of generating an RSA key pair with an OpenSSH UNIX client. In the example, the `identity.pub` file is the public-key file that you copy to the storage system root volume.

```
% ssh-keygen -t rsa -b 1024
Generating public/private rsa 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
```

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.

About this task

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/ssh/user_name/.ssh/authorized_keys2` file.

Examples 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/sshhd/user_name/.ssh/authorized_keys2` file on storage system `sys1`:

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

Related tasks

[Editing public keys generated by SecureCRT and ssh.com clients](#) on page 43

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.

Examples 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"
AAAAB3NzaC1yc2EAAAADAQABAAQgQDhJ6nk+2hm5iZnx737ZqxfgksPl3+OY1cP80s
```

```
1amXuUrwBp3/MUODEP5E511zqjO0w5kyJlvPjCiLg9UqS7JeY5yd/6xyGarsde26De1E
rbVJluqnxyAO1V9AlhjBE8TbI+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 AAAAB3NzaC1yc2EAAAADAQABAAQGDJhJ6nk+2hm5iZnx737ZqxfgksPl
3+OY1cP8s1amXuUrwBp3/MUODEP5E511zqjO0w5kyJlvPjCiLg9UqS7JeY5yd/6xy
Garsde26De1ErbVJluqnxyAO1V9AlhjBE8TbI+lyYBH+WezT0nySix6VBQTAWhv43r
9lSudswYV80Q==
```

Issuing SSH requests

You can issue SSH requests to the storage system to perform administrative tasks.

Before you begin

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, SecureAdmin must have been set up and enabled.

About this task

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] [-6] username@{IP_addr|hostname} [command]
```

or

```
ssh [-1|-2] [-6] -l username {IP_addr|hostname} [command]
```

- The option `-1` forces SSH to use protocol version 1 only.
SSH protocol version 1 supports only IPv4 addresses.
- The option `-2` forces SSH to use protocol version 2 only.
By default, SSH uses protocol version 2.
- The option `-6` is supported only for SSH protocol version 2 and forces SSH to use IPv6 addresses only.

Data ONTAP supports IPv4 addresses. If you use SSH protocol version 2 to access the storage system, and if options `ip.v6.enable` is set to on, IPv6 addresses are also supported.

For information about how to configure your system to use IPv6 addresses, see the *Data ONTAP Software Setup Guide for 7-Mode*.

- `command` is not required for SSH-interactive sessions.

Examples of SSH requests

The following examples show how the user account named “joe” can issue an SSH request to access the storage system:

```
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
```

```
ssh -2 joe@3FFE:81D0:107:2082::33 version
```

```
ssh -2 -6 joe@mysystem
```

In addition, if you use SSH protocol version 2 and if options `ip.v6.enable` is set to on, you can also specify IPv6 address information in the options `ssh.access` command, as shown in the following examples:

```
options ssh.access host=mysystem,10.72.137.28,3FFE:81D0:107:2082::33
```

```
options ssh.access "host = 3FFE:81D0:107:2082::33"
```

Related concepts

[How to manage SSH](#) on page 38

[Rules that apply to console, Telnet, and SSH-interactive sessions](#) on page 29

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

```
mssystem> options ssh
ssh.access                legacy
ssh.enable                on
ssh.idle.timeout          0
ssh.passwd_auth.enable    on
ssh.port                  22
ssh.pubkey_auth.enable    on
ssh1.enable               on
ssh2.enable               on
mssystem>
```

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.

Using SSL improves security by encrypting the administrator's password and all administrative communication when you manage your system from a browser.

Data ONTAP supports SSLv2, SSLv3, and Transport Layer Security version 1.0 (TLSv1.0). You should use TLSv1.0 or 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.

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 (CA) signed certificate**
A CA-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.
To enhance security, starting with Data ONTAP 8.0.2, Data ONTAP uses the SHA256 message-digest algorithm to generate digital certificates (including CSRs and root certificates) on the storage system.

Related concepts

[The default security settings](#) on page 35

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 CA signed certificate](#) on page 48

Installing a CA signed certificate

The advantage of a Certificate Authority (CA) 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.

The `secureadmin.pem` file is in the `/etc/keymgr/csr` directory on the storage system.

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. Reenable SSL by entering the following command:

```
secureadmin enable ssl
```


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.

About this task

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 47

Enabling or disabling SSL

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

Before you begin

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 47

[Installing a CA signed certificate](#) on page 48

Enabling or disabling SSLv2 or SSLv3

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

About this task

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.

TLS offers better security than SSLv3, and SSLv3 offers better security than SSLv2. In addition to enabling the SSL protocol, you must also have at least one of SSLv2, SSLv3, or TLS 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, FTPS, and LDAP connections, if the following options are also set to `on`:

- `httpd.admin.ssl.enable` (for HTTPS)
- `ftpd.implicit.enable` or `ftpd.explicit.enable` (for FTPS)
- `ldap.ssl.enable` (for LDAP)

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

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

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

Related tasks

- [Setting up and starting SSL](#) on page 47
- [Enabling or disabling TLS](#) on page 51

Enabling or disabling TLS

Enabling Transport Layer Security (TLS) enables the storage system to use TLS on HTTPS, FTPS, and LDAP traffic.

Before you begin

TLS is disabled by default, and setting up SSL does not automatically enable TLS. Before enabling TLS, ensure that SSL has been set up and enabled.

About this task

Data ONTAP supports TLSv1, SSLv3, and SSLv2. TLSv1 is a protocol version higher than SSLv3, and SSLv3 is a protocol version higher than SSLv2. A negotiation process is built into the TLS and the SSL protocols to use the highest protocol version that is supported by both the client and the server for communication. For TLS to be used for communication, both the client requesting connection and the storage system must support TLS.

Step

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

```
options tls.enable {on|off}
```

- Use `on` to enable TLS.
 - For TLS to take effect on HTTPS, ensure that the `httpd.admin.ssl.enable` option is also set to `on`.
 - For TLS to take effect on FTPS, ensure that the `ftpd.implicit.enable` option or the `ftpd.explicit.enable` option is also set to `on`.
 - For TLS to take effect on LDAP, ensure that the `ldap.ssl.enable` option is also set to `on`.

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

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

- Use `off` (the default) to disable TLS.

When TLS is disabled, SSL is used for communication if SSL has previously been set up and enabled.

Related tasks

[Determining whether secure protocols are enabled](#) on page 52

[Setting up and starting SSL](#) on page 47

[Installing a CA signed certificate](#) on page 48

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 35

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.

Related concepts

The default security settings on page 35

Rules that apply to console, Telnet, and SSH-interactive sessions on page 29

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 Enter 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 29

Related tasks

[Restricting protocol access](#) on page 65

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.

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 that 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.

Before you begin

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.

Related concepts

[*The default security settings*](#) on page 35

[*How to specify administration hosts*](#) on page 63

[*Public-key-based authentication*](#) on page 41

Related tasks

[*Restricting protocol access*](#) on page 65

[*Adding administration hosts*](#) on page 64

[*Removing administration hosts*](#) on page 64

[*Restricting protocol access*](#) on page 65

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

The `rsh.enable` option must be set to on.

If you access the storage system by using its IPv6 address, the `ip.v6.enable` option must be set to on for the system and the UNIX client you use must support IPv6.

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, IPv4 address, or IPv6 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 uses a user name, `carl`, and a password, `mypass`, to access the storage system whose IPv6 address is `2001:0DB8:85A3:0:0:8A2E:0370:99` to run the Data ONTAP `version` command:

```
rsh 2001:0DB8:85A3:0:0:8A2E:0370:99 -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 65

Accessing a storage system from a Windows client by using RSH

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

Before you begin

The `rsh.enable` option must be set to `on`.

The Windows client you use must be a trusted host specified in the `/etc/hosts.equiv` file on the root volume of the storage system.

If you access the storage system by using its IPv6 address, the `ip.v6.enable` option must be set to `on` for the system and the Windows client you use must support IPv6.

Step

1. From the RSH application on the Windows client, enter the `rsh` command in the following format:

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

hostname_or_ip is the host name, IPv4 address, or IPv6 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 uses a user name, “carl”, and a password, “mypass”, to access the storage system whose IPv6 address is 2001:0DB8:85A3:0:0:8A2E:0370:99, to run the Data ONTAP `version` command:

```
rsh 2001:0DB8:85A3:0:0:8A2E:0370:99 -l carl:mypass version
```

Related tasks

[Restricting protocol access](#) on page 65

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 IPv4 or IPv6 address for the RSH session

Note: If the `ip.v6.enable` option is set to `off`, `rshstat -a` displays only IPv4 connections.

- 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
```

Understanding OnCommand System Manager

System Manager is a graphical management interface that enables you to manage storage systems and storage objects (such as disks, volumes, and aggregates) and perform common management tasks related to storage systems from a Web browser.

You can use System Manager to manage storage systems running the following versions of Data ONTAP:

- Data ONTAP 7.3.x (starting from 7.3.7)
- Data ONTAP 8.0.4 and 8.0.5 operating in 7-Mode

- Data ONTAP 8.1.2
- Data ONTAP 8.2

You can also use System Manager to manage V-Series systems.

System Manager enables you to perform many common tasks such as the following:

- Configure and manage storage objects, such as disks, aggregates, volumes, qtrees, and quotas.
- Configure protocols, such as CIFS and NFS, and provision file sharing.
- Configure protocols such as FC and iSCSI for block access.
- Verify and configure network configuration settings in the storage systems.
- Create and manage vFiler units.
- Set up and manage SnapMirror relationships and SnapVault relationships.
- Manage HA configurations and perform takeover and giveback operations.

Note: System Manager replaces FilerView as the tool to manage storage systems running Data ONTAP 8.1 or later.

For more information about System Manager, see the NetApp Support Site.

Related information

NetApp Support Site: support.netapp.com

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.

Reasons to designate a workstation as an administration 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, and 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 Software Setup Guide for 7-Mode*.

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 privileges enough 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 privileges enough 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
2001:0DB8:85A3:0:0:8A2E:0370:99 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.

Related concepts

The default security settings on page 35

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>

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

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*.
The **ssh.access** and **rsh.access** options support both IPv4 and IPv6 addressing.
- *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 Network Management Guide for 7-Mode*.

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

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

Related tasks

[Allowing only secure access to the storage system](#) on page 69

Controlling the NFS 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 wafl.root_only_chown on</code>
You want to allow the privilege of changing directory and file ownership to all users	Enter the following command: <code>options wafl.root_only_chown off</code>

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 you...	Enter the following command...
Do not want to set access restrictions for anonymous share lookups	<code>options cifs.restrict_anonymous 0</code>
Do not want to allow enumeration of users and shares	<code>options cifs.restrict_anonymous 1</code>
Want to fully restrict anonymous share lookups	<code>options cifs.restrict_anonymous 2</code>

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 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>
<code>telnet.access</code>	<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>
<code>telnet.distinct.enable</code>	<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>

Option	Description
<code>rsh.access</code>	Controls which hosts can access the storage system through a Remote Shell session for administrative purposes. You can restrict Remote Shell access to the storage system by specifying host names, IP addresses, or network interface names.
<code>ssh.access</code>	Controls which hosts can access the storage system through a Secure Shell session for administrative purposes. You can restrict Secure Shell access to the storage system by specifying host names, IP addresses, or network interface names.
<code>nfs.mount_rootonly</code>	Controls whether the storage system's volumes can be mounted from NFS clients only by the root user on privileged ports (ports 1 through 1,023) or by all users on all ports. This option is applicable only if the NFS protocol is licensed.
<code>wapl.root_only_chown</code>	Controls whether all users or only the root user can change directory and file ownership. This option is applicable only if the NFS protocol is licensed.
<code>cifs.restrict_anonymous</code>	Controls whether anonymous CIFS users can look up CIFS shares, users, or groups on a storage system. This option is applicable only if the CIFS protocol is licensed.

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 65

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.

If you want to enable FTPS and SFTP, see the *Data ONTAP File Access and Protocols Management Guide for 7-Mode*.

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.
SSHv1	<code>options ssh1.enable off</code> Note: Ensure that the <code>ssh.enable</code> option and the <code>ssh2.enable</code> option are set to on.

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 35

[The default security settings](#) on page 35

[Data ONTAP options for managing password rules](#) on page 124

Managing the root aggregate and the root volume

The root aggregate contains the root volume. The root volume contains special directories and configuration files for the storage system. You can access the directories within the root volume from an NFS or a CIFS client or by using FTP or SFTP. You can also designate a different volume to be the new root volume.

Understanding the root aggregate

The root aggregate contains the root volume, which contains special directories and configuration files that help you administer the storage system.

The following facts apply to the root aggregate:

- Starting with Data ONTAP 8.1, new systems are shipped with the root volume in a 64-bit root aggregate.
- 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 is available on the system.

- A Flash Pool (an aggregate that contains both HDDs and SSDs) can be used as the root aggregate.

Caution: If you revert or downgrade to Data ONTAP 8.1 or earlier with a Flash Pool configured as your root aggregate, your system will not boot.

Understanding the root volume

The storage system's root volume contains special directories and configuration files that help you administer the storage system. Understanding the facts about the root volume helps you manage it.

The following facts apply to the root volume:

- How the root volume is installed and whether you need to create it yourself depend on the storage system:
 - For FAS systems and V-Series systems ordered with disk shelves, the root volume is a FlexVol volume that is installed at the factory.
 - For a V-Series system that does not have a disk shelf, you install the root volume on an array LUN.

For more information about setting up a V-Series system, see the *Data ONTAP Software Setup Guide for 7-Mode*.

- For systems running virtual storage, the Data ONTAP-v installation process creates a single aggregate by using all currently defined virtual disks and creates the root FlexVol volume in that aggregate.
For more information about system setup, see the Installation and Administration Guide that came with your Data ONTAP-v system.
- The default name for the root volume is `/vol/vol0`.
You can designate a different volume to be the new root volume. Starting in Data ONTAP 8.0.1, you can designate a 64-bit volume to be the new root volume.
- The root volume's fractional reserve must be 100%.

Recommendations for the root volume

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

The following are general recommendations for root volumes:

- Root volumes can use either FlexVol or traditional volumes.
If a root volume exists as a traditional volume, it can be a stand-alone RAID4 or RAID-DP volume. RAID4 requires a minimum of two disks and can protect against single-disk failures. RAID-DP, the default RAID type, requires a minimum of three disks and can protect against double-disk failures. Using RAID-DP for the root aggregate is recommended.
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 information about RAID types, traditional volumes, and FlexVol volumes, see the *Data ONTAP Storage Management Guide for 7-Mode*.
- It is recommended that the root volume be in a separate aggregate that does not include data volumes or other user data.
However, for small storage systems where cost concerns outweigh resiliency, a FlexVol based root volume on a regular aggregate might be more appropriate.
- You should 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:

- 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

[Sizing considerations for root FlexVol volumes](#) on page 73

Related tasks

[Changing the root volume](#) on page 83

Sizing considerations for root FlexVol volumes

The root volume must have enough space to contain system files, log files, and core files. It is important to meet the size requirements because, if a system problem occurs, these files are needed to provide technical support.

The following considerations apply to size requirements for root FlexVol volumes:

- 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.

For information about the minimum size for the root FlexVol volume for your storage system model, see the *Hardware Universe* (formerly the *System Configuration Guide*) at support.netapp.com/knowledge/docs/hardware/NetApp/syscfg/index.shtml.

- Data ONTAP prevents you from resizing the root volume below the minimum allowed size or changing the guarantee for the root volume.
If you designate a different volume to be the new root volume, the volume to be used as the new root volume must meet the minimum size requirement.
- If you are using third-party storage, the array LUN that is used for the root volume must be large enough to meet the minimum size requirement for the root volume.
See the *Hardware Universe* (formerly the *System Configuration Guide*) at support.netapp.com/knowledge/docs/hardware/NetApp/syscfg/index.shtml for the minimum array LUN size for the root volume, which is larger than the minimum array LUN size for data LUNs.
- The root Flexvol volume must not be increased to more than 95 percent of the available aggregate size.
The output of the `df -A` command displays the space used by the aggregates in the system.

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.

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.

What the configuration files are

Configuration files affect the behavior of the storage system. You can use an editor on your administration host to edit some configuration files in the `/etc` directory.

If a configuration file can be edited by the system administrator, it is listed in Section 5 of the man pages.

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

Related concepts

[Startup configuration for the storage system](#) on page 146

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.
/etc/netgroup	Maximum entry size of 4,096 characters. Maximum netgroup nesting limit is 1,000. No file size limit. Netgroup lookup is case-sensitive and must match the case used by DNS or NIS servers for host lookup.

File name	Limits
<code>/etc/passwd</code>	Maximum line size of 256 characters. No file size limit.
<code>/etc/resolv.conf</code>	Maximum line size is 256. Maximum number of name servers is 3. Maximum domain name length is 256. Maximum search domains limit is 6. Total number of characters for all search domains is limited to 256. No file size limit.

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.

About this task

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.

About this task

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

[Understanding message logging](#) on page 142

[How to access the default directories on the storage system](#) on page 79

Related tasks

[Accessing log files using HTTP or HTTPS](#) on page 82

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.

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 File Access and Protocols Management Guide for 7-Mode* 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 35

Accessing the /etc directory with SFTP

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

Before you begin

SFTP requires SSHv2. Before enabling SFTP, ensure that SSHv2 has been set up and enabled for your storage system.

Steps

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

```
options sftp.enable on
```

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

```
options sftp.dir_override /vol/vol0/etc
```

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

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

Related tasks

[Determining whether secure protocols are enabled](#) on page 52

[Displaying the current SSH settings](#) on page 46

[Setting up and starting SSH](#) on page 39

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 File Access and Protocols Management Guide for 7-Mode* 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 35

Accessing the /home directory with SFTP

You can use the SSH File Transfer Protocol (SFTP) to access the /home directory of your storage system.

Before you begin

SFTP requires SSHv2. Before enabling SFTP, ensure that SSHv2 has been set up and enabled for your storage system.

Steps

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

```
options sftp.enable on
```

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

```
options sftp.dir_override /vol/vol0/home
```

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

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

Related tasks

[Determining whether secure protocols are enabled](#) on page 52

[Displaying the current SSH settings](#) on page 46

[Setting up and starting SSH](#) on page 39

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 35

Related tasks

Allowing only secure access to the storage system on page 69

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 is used as the system's root volume.

Before you begin

The volume that you are designating to be the new root volume must meet 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.

In addition, the volume that you are designating to be the new root volume must have at least 2 GB of free space. It must also have a fractional reserve of 100%. The `vol status -v` command displays information about a volume's fractional reserve.

If you use a FlexVol volume for the root volume, ensure that it has a guarantee of `volume`. For information about volume guarantees, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Starting in Data ONTAP 8.0.1, you can designate a volume in a 64-bit aggregate to be the new root volume.

If you move the root volume outside the current root aggregate, you must also change the value of the aggregate `root` option so that the aggregate containing the root volume becomes the root aggregate.

For V-Series systems with the root volume on the storage array, the array LUN used for the root volume must meet the minimum array LUN size for the root volume. For more information about the minimum array LUN size for the root volume on V-Series systems, see the *Hardware Universe* (formerly the *System Configuration Guide*) at support.netapp.com/knowledge/docs/hardware/NetApp/syscfg/index.shtml.

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. For information about changing your root volume from a traditional volume to a FlexVol volume, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Steps

1. Identify an existing volume to use as the new root volume, or create the new root volume by using the `vol create` command.

For more information about creating volumes, see the *Data ONTAP Storage Management Guide for 7-Mode*.

2. Use the `ndmcopy` command to copy the `/etc` directory and all of its subdirectories from the current root volume to the new root volume.

For more information about `ndmcopy`, see the *Data ONTAP Data Protection Tape Backup and Recovery Guide for 7-Mode*.

3. Enter the following command to specify the new root volume:

```
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.

4. If you moved the root volume outside the current root aggregate, enter the following command to change the value of the aggregate `root` option so that the aggregate containing the root volume becomes the root aggregate:

```
aggr options aggr_name root
```

`aggr_name` is the name of the new root aggregate.

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

5. 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.

If you changed the root aggregate, a new root volume is created during the reboot when the aggregate does not already contain a FlexVol volume designated as the root volume and when the aggregate has at least 2 GB of free space.

6. Update the `httpd.rootdir` option to point to the new root volume.

Related concepts

Recommendations for the root volume on page 72

Sizing considerations for root FlexVol volumes on page 73

Starting and stopping the storage system

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

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 Upgrade and Revert/Downgrade Guide for 7-Mode*.

Ways to boot the storage system

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

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>` or `RLM>`.

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.

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.

Rebooting the storage system at the system prompt

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

Step

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

```
reboot
```

The storage system begins the reboot process. The Data ONTAP login prompt appears, indicating that the reboot process is complete.

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 to access the boot environment prompt:

```
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	boot_ontap
The Data ONTAP primary image from the boot device	boot_primary
The Data ONTAP backup image from the boot device	boot_backup

Rebooting the storage system remotely

You can reboot the 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
```

One of the following remote management device CLI prompts appears, depending on the platform model:

```
SP>  
RLM>
```

2. To power 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
```

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

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

Related concepts

[Ways to boot the storage system](#) on page 86
[Managing a storage system remotely](#) on page 189

Recovering from a corrupted image of the system's 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 appears. It can be one of the following, depending on the platform model:

```
SP>
```

```
RLM>
```

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 BIOS firmware image, enter the following command:

```
system reset primary
```

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

```
system reset backup
```

The console informs you that the command will cause a “dirty system shutdown” and asks you whether to continue.

3. Enter **y** to continue.

The storage system shuts down immediately.

Related concepts

[Ways to boot the storage system](#) on page 86

[Managing a storage system remotely](#) on page 189

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	version
Data ONTAP versions available on the boot device	version -b

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.

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
```

One of the following remote management device CLI prompts appears, depending on the storage system model:

```
SP>
```

```
RLM>
```

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 86

[Managing a storage system remotely](#) on page 189

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.

Managing the storage system by using the boot menu

You can use the boot menu to correct configuration problems, reset the root password, initialize disks, reset system configuration, and restore system configuration information back to the boot device.

Steps

1. Reboot the system to access the boot menu by entering the following command at the system prompt:

reboot

The storage system begins the reboot process.

2. During the reboot process, press Ctrl-C to display the boot menu when prompted to do so.

The storage system displays the following options for the 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)?
```

3. Select one of the following options by entering the corresponding number:

To...	Select...
Continue to boot the storage system in normal mode	1) Normal Boot

To...	Select...
Troubleshoot and repair configuration problems	2) Boot without /etc/rc This menu option causes the storage system to do the following: <ul style="list-style-type: none"> • Use only default options settings • Disregard all options settings you specified in the <code>/etc/rc</code> file • Disable some services, such as <code>syslog</code>
Change the password of the storage system	3) Change Password
Initialize all the disks and create a root FlexVol volume	4) Clean configuration and initialize all disks <p>Attention: This menu option erases all data on the disks and resets your system configuration to the factory default settings. If you need to preserve existing configuration values that are used for system setup (such as your system IP address, gateway addresses, and DNS server addresses), you should make a note of the values before selecting this menu option. You can find your current setup settings by using the <code>setup</code> command at the system prompt.</p> <p>This menu option reboots the storage system before initializing the disks. After the initialization procedure finishes, the setup script starts and prompts you for configuration information. For information about setting up the storage system, see the <i>Data ONTAP Software Setup Guide for 7-Mode</i>. Data ONTAP 8.0 and later releases do not allow you to create a new traditional root volume from the boot menu. However, preexisting traditional root volumes are still supported.</p> <p>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>
Perform aggregate and disk maintenance operations and obtain detailed aggregate and disk information.	5) Maintenance mode boot You exit Maintenance mode by using the <code>halt</code> command.
Restore the configuration information from the root FlexVol volume to the boot device, such as a PC CompactFlash card	6) Update flash from backup config Data ONTAP stores some system configuration information on the boot device. When the storage system reboots, the information on the boot device is automatically backed up onto the root FlexVol 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 FlexVol volume back to the boot device.

To...	Select...
Install new software on a V-Series system	7) Install new software first If the Data ONTAP software on the boot device does not include support for the storage array that 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. 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.
Reboot the storage system	8) Reboot node

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

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.

Reasons for creating administrator accounts

In addition to using the default system administration account (“root”) for managing a storage system, you can create additional administrator user accounts to manage administrative access to the storage system.

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 enables you to display information about who is performing which 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, domain user accounts authenticated on the Windows domain can access the storage system using Telnet, RSH, SSH, 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 File Access and Protocols Management Guide for 7-Mode*.

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.

domain user: 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 domain users 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. Examples of types of capabilities include the following:
 - 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 domain users. 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 domain users 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.

Related concepts

Supported capability types on page 108

Disabling root account 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.

Before you begin

You can disable the root account's access to the storage system if all the following conditions are met:

- You user account is not “root”.
- You have the `security-complete-user-control` security capability.
- You use Telnet, RSH, SSH, HTTP Admin, NDMP, or the serial console to access the storage system.

About this task

The storage system's root account is mapped to the “naroot” account name of the remote management device. If you disable the root account's access to the storage system, the naroot access to the storage system is automatically disabled.

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.

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]
```

- You use `useradmin user add` to create a new user, and `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 quotation marks (" ").

- *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 creating a user

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 103

[Requirements for naming users, groups, and roles](#) on page 96

Related tasks

[Assigning roles to groups by creating or modifying a group](#) on page 104

Granting access to Windows domain users

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

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.

Data ONTAP also supports ssh-key based authentication for domain users.

Steps

1. To assign a Windows domain user 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 domain user whose name or Security ID (SID) you want to assign to a customized or predefined group. This value can be in one of the following formats:

- `name`

Note: If you do not specify the domain name, the domain is the storage system, and the user is considered distinct from any user in the Windows domain with the same user name.

- `domain\name`
- `textual_sid_S-x-y-z`

For more information about these formats, see the `na_cifs_lookup(1)` man page.

`custom_group` is a customized group with roles assigned through the `useradmin group` command.

`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 99

[Predefined groups](#) on page 103

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 99

[Granting access to Windows domain users](#) on page 101

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 aging 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.

Predefined groups

You can assign a user or domain user 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 off.</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 106[Supported capability types](#) on page 108**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 96[Windows special groups](#) on page 97

[Predefined roles](#) on page 106

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 97

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.

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.

Predefined roles

Data ONTAP provides predefined roles that you can assign to groups and enable users in the groups to perform different levels of administrative tasks.

The following table describes the roles that are predefined by Data ONTAP.

This role...	Includes the following default capabilities...	That grant users of a group...
root	*	All possible capabilities.
admin	cli-*, api-*, login-*, security-*	All CLI, API, login, and security capabilities.

This role...	Includes the following default capabilities...	That grant users of a group...
power	cli-cifs*, cli-exportfs*, cli-nfs*, cli-useradmin*, api-cifs-*, api-nfs-*, login-telnet, login-http-admin, login-rsh, login-ssh, api-system-api-*	<p>The capabilities to performing the following tasks:</p> <ul style="list-style-type: none"> • Invoke all cifs, exportfs, nfs, and useradmin CLI commands • Make all cifs and nfs API calls • Log in using Telnet, HTTP, RSH, and SSH sessions
backup	login-ndmp	The capabilities to make NDMP requests.
compliance	cli-cifs*, cli-exportfs*, cli-nfs*, cli-useradmin*, api-cifs-*, api-nfs-*, login-telnet, login-http-admin, login-rsh, login-ssh, api-system-api-*, cli-snaplock*, api-snaplock-*, api-file-*, compliance-*	<p>Compliance-related capabilities in addition to all the capabilities granted by the power role.</p> <p>Note: The compliance role is the default role for the Compliance Administrators group. The compliance role cannot be removed from the Compliance Administrators group or added to other groups.</p>
audit	api-snmp-get, api-snmp-get-next	The capabilities to make snmp-get and snmp-get-next API calls.
none	None	No administrative capabilities.

Related concepts

[Predefined groups](#) on page 103

[Supported capability types](#) on page 108

Related tasks

[Assigning roles to groups by creating or modifying a group](#) on page 104

Supported capability types

The capability types Data ONTAP supports include `login`, `cli`, `security`, `api`, and `compliance`.

The following table describes the supported capability types.

This capability type...	Has the following capabilities...
<code>login</code>	<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.
<code>cli</code>	<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...
security	<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</code> advanced 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.
api	<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>

This capability type...	Has the following capabilities...
compliance	<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>

Related concepts

[About changing another user's capabilities](#) on page 102

[Predefined roles](#) on page 106

[Predefined groups](#) on page 103

Related tasks

[Loading groups from the `lclgroups.cfg` file](#) on page 105

[Creating a new role and assigning capabilities to roles](#) on page 110

[Assigning roles to groups by creating or modifying a group](#) on page 104

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 102

[Requirements for naming users, groups, and roles](#) on page 96

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.

Commands that list users, domain users, groups, or roles

You use the `useradmin` commands to display information for users, domain users, 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 <i>user_name</i></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 <i>grp_name</i></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 Security IDs (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 <code>useradmin user list</code> command output includes the Relative ID (RID), which is a unique integer associated with each user. 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, domain users, groups, or roles

You use the `useradmin` commands to delete users, domain users, 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>
<code>useradmin domainuser delete win_user_name -g group1,[group2,...]</code>	<p>Removes the specified user from the specified group or groups. User names are case insensitive.</p> <p>This command does not delete the user from the domain.</p> <p>Note: If you want to completely delete a user from the storage system, use the <code>useradmin user delete</code> command instead.</p>
<code>useradmin group delete group_name</code>	<p>Deletes the specified group from the storage system. Group names are case insensitive.</p> <p>Note: All users must be removed from a group before the group itself can be deleted.</p>

Command	Description
<code>useradmin role delete <i>role_name</i></code>	<p>Deletes the specified role from the storage system. Role names are case insensitive.</p> <p>Note: A role that is still assigned to a group cannot be deleted.</p>

Administrative user creation examples

You can create a user with custom capabilities or no administrative capabilities, thereby controlling the user's administrative access.

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”, which is allowed to log in only via SSH
 - “qtree_commands”, which can run any `qtree` command in the CLI
- Create the “ssh_qtree_admins” group, which is allowed to log in only via SSH and run the `qtree` commands in the CLI, by using the `only_ssh` and `qtree_commands` roles.
- Create a user “wilma” and assign the 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.

Granting users in LDAP groups access to the system and mapping them to specified roles

If you store your user database on an LDAP server, you can grant users in LDAP groups access to the storage system and map them to specified roles on the system to manage their access.

Steps

1. If the value of the `security.admin.authentication` option does not include `nsswitch`, add `nsswitch` to the option by using one of the following formats:

- **`options security.admin.authentication internal,nsswitch`**
- **`options security.admin.authentication nsswitch,internal`**

The `security.admin.authentication` option specifies where the system finds authentication information for administrative user accounts. By default, it includes `internal`, which means the system's local administrative repository. Adding `nsswitch` to the option enables the system to also use the repositories found in the `nsswitch.conf` file.

For more information about the `security.admin.authentication` option, see the `na_options(1)` man page. For information about configuring LDAP services and the `nsswitch.conf` file, see the *Data ONTAP File Access and Protocols Management Guide for 7-Mode*.

2. To grant users in LDAP groups access to the storage system and map them to specified roles on the system, enter the following command:

```
options security.admin.nsswitchgroup  
ldapgroup1:role1,ldapgroup2:role2,ldapgroup3:role3...
```

- The `security.admin.nsswitchgroup` option maps an LDAP group to the role that follows the colon (:) after the group name.
For instance, `ldapgroup1` is mapped to `role1`, `ldapgroup2` to `role2`, and `ldapgroup3` to `role3`.
- Group names and role names must not contain commas (,) or colons (:).
- Mapping an LDAP group to a role enables users in that group to have only the capabilities granted by the mapped role.
The role can be a predefined role or one that you create by using the `useradmin role add` command.
- If you use the option without specifying a role after an LDAP group, users in that group are granted capabilities of the `admin` role and have full administrative access to the storage system.
- The `security.admin.nsswitchgroup` option supports up to 256 characters and ignores characters that exceed the length limit.
For more information about the option, see the `na_options(1)` man page.

Example

The following example grants LDAP users in the `ldapgrp1` group capabilities defined in the `power` role, LDAP users in the `ldapgrp2` group full administrative capabilities, and LDAP users in the `ldapgrp3` group capabilities defined in the `audit` role:

```
system> options security.admin.nsswitchgroup ldapgrp1:power,ldapgrp2,  
ldapgrp3:audit
```

Related concepts

[How to manage roles](#) on page 106

How to manage passwords for security

Data ONTAP provides several methods that you can use to ensure that the password policies for your storage system meet your company's security requirements.

When the `security.passwd.rules.enable` option is set to on (the default), you can manage passwords in the following ways:

- Password composition rules

The following are the default password composition rules for all accounts, including the “root” and “Administrator” accounts:

- The password must be at least eight characters long.
The `security.passwd.rules.minimum` option defaults to 8.
- The password must contain at least one number.
The `security.passwd.rules.minimum.digit` option defaults to 1.
- The password must contain at least two alphabetic characters.
The `security.passwd.rules.minimum.alphabetic` option defaults to 2.
- The password must not contain the Ctrl-c or Ctrl-d key combination or the two-character string `^D`.

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 Software Setup Guide for 7-Mode*. Also, even if you modified the settings for the password composition rules, the default settings still apply when you use the boot menu option **3**)

Change Password to change the system password (the “root” account password.)

By default, a password is not required to include symbol characters, but you can change the requirement by using the `security.passwd.rules.minimum.symbol` option.

In addition, you can use the following options to specify the minimum number of uppercase or lowercase alphabetic characters that a password must contain:

- The `security.passwd.rules.minimum.uppercase` option specifies the minimum number of uppercase alphabetic characters that a password must contain.
The default is 0, which does not require that a password contain uppercase characters.
- The `security.passwd.rules.minimum.lowercase` option specifies the minimum number of lowercase alphabetic characters that a password must contain.
The default is 0, which does not require that a password contain lowercase characters.
- Password history

The password history functionality 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 must 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 most recent 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.

- Password expiration (maximum age)

The password expiration functionality 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.

Note: Before you set password expiration, your storage system time must already be set accurately. Otherwise, accounts could expire before or after the set expiration date. The settings of the `useradmin` commands do not depend on the `security.passwd.rules.enable` option. They do not require that `security.passwd.rules.enable` be set to on.

- Password minimum age

The password minimum age functionality (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.

Note: Before you set the password minimum age, your storage system time must already be set accurately. Changing the system time after the password minimum age is set can lead to unexpected results.

- Password lockout

The password lockout functionality 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.

Note: The setting of the `security.passwd.lockout.numtries` option does not depend on the `security.passwd.rules.enable` option. It does not require that `security.passwd.rules.enable` be set to on.

- 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.

Note: The setting of the `security.passwd.firstlogin.enable` option does not depend on the `security.passwd.rules.enable` option. It does not require that `security.passwd.rules.enable` be set to on.

For more information about options that manage passwords, see the `na_options(1)` and `na_useradmin(1)` man pages.

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" account name, 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 35

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	<ol style="list-style-type: none">a. Enter the following command: passwdb. 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	<p>Enter the following command:</p> <pre>rsh system_name -l username:password passwd old_password new_password username</pre>
Secure Shell connection	<p>Enter the following command:</p> <pre>ssh -l username system_name passwd old_password new_password username</pre>

Related concepts

[The default security settings](#) on page 35

Data ONTAP options for managing password rules

Data ONTAP provides several options that you can use to manage password rules. You can specify password requirements such as how a check for password composition is performed and what the maximum or minimum number of characters a password requires.

The following Data ONTAP options enable you to manage password rules:

This option (used with the options command)...	Enables you to...
<code>security.passwd.firstlogin.enable</code>	<p>Specify 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 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>
<code>security.passwd.lockout.numtries</code>	<p>Specify the number of allowed login attempts before a nonroot user's account is disabled.</p> <p>The default value is 4,294,967,295.</p>
<code>security.passwd.rules.enable</code>	<p>Specify 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 with options that begin with <code>security.passwd.rules</code>, and a password is rejected if it does not pass the check. If this option is set to <code>off</code>, the check is not performed.</p> <p>The default value is <code>on</code>.</p> <p>This option does not apply to the users “root” or “Administrator” (the NT Administrator account) if <code>security.passwd.rules.everyone</code> is set to <code>off</code>.</p>
<code>security.passwd.rules.everyone</code>	<p>Specify whether a check for password composition is performed for all users, including the users “root” and “Administrator”.</p> <p>If this option is set to <code>off</code>, the checks do not apply to “root” or “Administrator”. 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 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>

This option (used with the options command)...	Enables you to...
<code>security.passwd.rules.history</code>	<p>Specify the number of previous passwords that are checked against a new password to prevent repeats.</p> <p>For storage systems shipped with Data ONTAP 8.0 or later, The default value 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>
<code>security.passwd.rules.maximum</code>	<p>Specify the maximum number of characters a password can contain.</p> <p>The value of this option must not be smaller than that of <code>security.passwd.rules.minimum</code>.</p> <p>The default value is 256.</p> <p>Note: This option can be set to a value greater than 16, but a maximum of 16 characters are used to match the password. The system ignores characters that are beyond the first 16 when checking the password against the composition rules.</p> <p>Users with passwords longer than 14 characters cannot log in through the Windows interfaces. Therefore, 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>
<code>security.passwd.rules.minimum</code>	<p>Specify the minimum number of characters a password must contain.</p> <p>The value of this option must not be greater than that of <code>security.passwd.rules.maximum</code>.</p> <p>The default value is 8.</p> <p>If the <code>security.passwd.rules.enable</code> option is set to <code>off</code>, this option is ignored.</p>

This option (used with the options command)...	Enables you to...
<code>security.passwd.rules. minimum.alphabetic</code>	<p>Specify the minimum number of alphabetic characters a password must contain.</p> <p>This number includes the required numbers of uppercase and lowercase characters, which you can set by using <code>security.passwd.rules.minimum.uppercase</code> and <code>security.passwd.rules.minimum.lowercase</code> respectively.</p> <p>The default value is 2.</p> <p>If the <code>security.passwd.rules.enable</code> option is set to <code>off</code>, this option is ignored.</p> <p>If this option has a value smaller than the combined value of <code>security.passwd.rules.minimum.uppercase</code> and <code>security.passwd.rules.minimum.lowercase</code>, the uppercase and lowercase rules determine the required number of alphabetic characters in a password. In the following examples, the system ensures that a password contains at least two uppercase and four lowercase alphabetic characters:</p> <pre>options security.passwd.rules.minimum.alphabetic 0 options security.passwd.rules.minimum.uppercase 2 options security.passwd.rules.minimum.lowercase 4</pre> <pre>options security.passwd.rules.minimum.alphabetic 3 options security.passwd.rules.minimum.uppercase 2 options security.passwd.rules.minimum.lowercase 4</pre>
<code>security.passwd.rules. minimum.digit</code>	<p>Specify the minimum number of digit characters a password must contain. These are numbers from 0 to 9.</p> <p>The default value is 1.</p> <p>If the <code>security.passwd.rules.enable</code> option is set to <code>off</code>, this option is ignored.</p>
<code>security.passwd.rules. minimum.lowercase</code>	<p>Specify the minimum number of lowercase alphabetic characters (“a” to “z”) that a password must contain.</p> <p>The default value is 0.</p> <p>If the <code>security.passwd.rules.enable</code> option is set to <code>off</code>, this option is ignored.</p>

This option (used with the options command)...	Enables you to...
<code>security.passwd.rules.minimum.symbol</code>	Specify the minimum number of symbol characters (including white space and punctuation characters) a password must contain. The default value is 0. If the <code>security.passwd.rules.enable</code> option is set to <code>off</code> , this option is ignored.
<code>security.passwd.rules.minimum.uppercase</code>	Specify the minimum number of uppercase alphabetic characters (“A” to “Z”) that a password must contain. The default value is 0. If the <code>security.passwd.rules.enable</code> option is set to <code>off</code> , this option is ignored.

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

Uses of the systemshell and the diagnostic account

A diagnostic account, named “diag”, is provided with your storage system. You can use the diag account to perform troubleshooting tasks in the systemshell. The diag account and the systemshell are intended only for low-level diagnostic purposes and should be used only with guidance from technical support.

The diag account is the only account that can be used to access the systemshell, through the advanced command `systemshell1`. The diag account is disabled by default. You must enable the account and set up its password before using it for the first time. Neither the diag account nor the systemshell is intended for general administrative purposes.

Enabling or 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 the 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	useradmin diaguser show By default, the diagnostic account is disabled. 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> .
Enable the diagnostic account.	useradmin diaguser unlock
Disable the diagnostic account	useradmin diaguser lock

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.

```
systemname*> useradmin diaguser show
Name: diag
Info: Account for access to systemshell
Locked: yes

systemname*> useradmin diaguser unlock

systemname*> 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 the 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.
- The password must not contain the Ctrl-c or Ctrl-d key combination or the two-character string ^D.

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.

```
systemname*> 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.

```
systemname*> systemshell
login: diag
Password:

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

systemname*>
```

The following example shows the screen output of the `systemshell` command when the diagnostic account is disabled.

```
systemname*> useradmin diaguser lock

systemname*> useradmin diaguser show
Name: diag
Info: Account for access to systemshell
Locked: yes

systemname*> systemshell
login: diag
Password:
Login incorrect
login: diag
Password:
Login incorrect
(CTRL-C)

systemname*>
```

Related tasks

[Enabling or disabling the diagnostic account](#) on page 128

[Setting the password for the diagnostic account](#) on page 129

General system maintenance

Certain tasks are required for your system to run properly. The tasks include managing aggregate Snapshot copies; managing licenses; managing the system date and time; managing core dump files; configuring message logging, audit logging, and system startup; and backing up and cloning system configuration.

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`

Managing aggregate Snapshot copies

An aggregate Snapshot copy is a point-in-time, read-only image of an aggregate. You use aggregate Snapshot copies when the contents of an entire aggregate need to be recorded.

An aggregate Snapshot copy is similar to a volume Snapshot copy, except that it captures the contents of the entire aggregate, rather than any particular volume. Also, you do not restore data directly from an aggregate Snapshot copy. To restore data, you use a volume Snapshot copy.

How you use aggregate Snapshot copies depends on whether you use the SyncMirror or MetroCluster functionality.

- If you use SyncMirror or MetroCluster, you must enable automatic aggregate Snapshot copy creation and keep your aggregate Snapshot reserve at 5 percent or higher.
If you use SyncMirror or MetroCluster 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 resynchronize the mirror later.
Also, 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 you do not use either SyncMirror or MetroCluster, you do not need to enable automatic aggregate Snapshot copy creation or reserve space for aggregate Snapshot copies.
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. If that is important to you, you

can ensure that automatic aggregate Snapshot copy creation is enabled. However, disabling automatic aggregate Snapshot copy creation and keeping your aggregate Snapshot reserve at 0 percent increases your storage utilization, because no disk space is reserved for aggregate Snapshot copies. Disabling automatic aggregate Snapshot copy creation and setting the aggregate Snapshot reserve to 0 percent does not affect normal operation, except for making more free space available for data.

For more information about Snapshot copies, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

Considerations for increasing the aggregate Snapshot reserve

You should consider increasing the aggregate Snapshot reserve under some circumstances.

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*.

As more and more data blocks in the aggregate are changed, the aggregate Snapshot reserve gradually becomes full. Data ONTAP automatically deletes the oldest aggregate Snapshot copies to recover space in the aggregate Snapshot reserve.

Newly created nonmirrored aggregates, including the root aggregate, have the aggregate Snapshot reserve set to 0 percent by default. A newly created mirrored aggregate's Snapshot reserve is set to 5 percent by default.

If you convert an existing, nonmirrored aggregate to a mirrored aggregate, the system attempts to increase the aggregate Snapshot reserve to 5 percent. If there is not enough space in the aggregate for the reserve increase, the operation to convert a nonmirrored aggregate to a mirrored aggregate fails.

You should consider increasing the aggregate Snapshot reserve in the following situations:

- You find that aggregate Snapshot copies are being created and deleted often enough to affect system performance.
- You need to complete a mirror resync operation when data is being written to an aggregate frequently.

In this case, the standard aggregate Snapshot reserve size of 5 percent might not be large enough to hold all the resynchronized Snapshot copies until the resync operation is complete.

- You do not want aggregate Snapshot copies deleted immediately.

To preserve aggregate Snapshot copies, you would increase the reserve sufficiently to keep the aggregate Snapshot copies available for overwrites until they are no longer needed and delete all other Snapshot copies.

If an aggregate has an aggregate Snapshot reserve of 0 percent, the operation to create aggregate Snapshot copies still succeeds if the aggregate has enough free space available.

Managing licenses

A license is a record of one or more software entitlements. Installing license keys, also known as *license codes*, enables you to use certain features or services on your storage system.

Data ONTAP feature licenses are issued as *packages*, each of which contains multiple features or a single feature. A package requires a license key, and installing the key enables you to access all features in the package. For information about the license packages, see the knowledgebase article [Data ONTAP 8.2 Licensing Overview and References](#) on the NetApp Support Site.

Starting with Data ONTAP 8.2, all license keys are 28 characters in length. Licenses installed prior to Data ONTAP 8.2 continue to work in Data ONTAP 8.2 and later releases. However, if you need to reinstall a license (for example, you deleted a previously installed license and want to reinstall it in Data ONTAP 8.2 or later, or you perform a controller replacement procedure for a system running Data ONTAP 8.2 or later), Data ONTAP requires that you enter the license key in the 28-character format.

You can find license keys for your initial or add-on software orders at the NetApp Support Site under **My Support > Software Licenses**. For instance, you can search with the serial number of a system to find all license keys associated with the system. If you cannot locate your license keys from the Software Licenses page, you should contact your sales or support representative.

Data ONTAP enables you to manage licenses in the following ways:

- Add one or more license keys (`license add`)
- Display information about installed licenses (`license show`)
- Delete a license from the storage system (`license delete`)

In an HA configuration, you must keep the licensing entitlement between the two nodes consistent. If a takeover occurs, the takeover node can provide only the functionality for the licenses installed on it. If the takeover node does not have a license that was being used by the partner node to serve data, your HA pair loses functionality after a takeover.

Some features require that you enable certain options instead of or in addition to installing a license key. For information, see the knowledgebase article [Data ONTAP 8.2 Licensing Overview and References](#) on the NetApp Support Site.

Related information

[NetApp Support Site: support.netapp.com](http://support.netapp.com)

License types

Understanding license types helps you manage the licenses in a storage system.

A package can have one or more of the following types of license installed in the storage system. The `license show` command displays the installed license type or types for a package.

- Standard license (`license`)

A standard license is issued for a storage system with a specific system serial number (also known as a *controller serial number*). It is valid only for the system that has the matching serial number.

Note: The `sysconfig` command displays the serial number of the system.

Data ONTAP 8.2 and later releases treat a license installed prior to Data ONTAP 8.2 as a standard license. Therefore, in Data ONTAP 8.2 and later releases, the system automatically has the standard license for the package that the previously licensed functionality is part of. The `license show` command with the `-legacy yes` parameter indicates such licenses.

- Site license (`site`)

A site license is not tied to a specific system serial number and can be installed on any storage system that is covered by the site license agreement. After you install a site license on a storage system, the `license show` command displays the site license under the system serial number of the storage system.

- Evaluation license (`demo`)

An evaluation license is a temporary license that expires after a certain period of time (indicated by the `license show` command). It enables you to try certain software functionality without purchasing an entitlement. It is not tied to a specific serial number of a system.

Commands for managing licenses

You use the `license` commands to manage licenses for the storage system.

If you want to...	Use this command...
Add one or more licenses	<code>license add</code>
Display information about installed licenses, for example: <ul style="list-style-type: none"> • License package name and description • License type (<code>site</code>, <code>license</code>, or <code>demo</code>) • Expiration date, if applicable • The system that a package is licensed for • Whether the license was installed prior to Data ONTAP 8.2 (<code>legacy</code>) • Customer ID 	<code>license show</code> <p>Note: Some information is displayed only when you use the <code>-instance</code> parameter.</p>
Delete the license of a package from the storage system	<code>license delete</code>

For more information, see the man pages for the `license` commands.

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

While SnapMirror is running, if you use the `date` or `rdate` command to set the system to an earlier date, 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 Public Services web at support.ntp.org.

The `timed` daemon operates independently on each node in an HA pair.

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.

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 136

Data ONTAP options for managing system time

You can use the `timed` options to enable time synchronization and specify the servers to use for time synchronization.

The following table describes the `timed` options:

This option...	Enables you to...
<code>timed.enable</code>	Enable time synchronization.
<code>timed.proto</code>	Specify the protocol used to synchronize the system time. This option supports only <code>ntp</code> .

This option...	Enables you to...
<code>timed.servers</code>	Specify up to five time servers used by the <code>timed</code> functionality.

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

Example of system time synchronization

The following example configures the system to use the NTP protocol.

```
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
```

Managing core dump files

When a hardware or software failure causes the storage system to panic, a core dump occurs and the system creates a core dump file that technical support can use to troubleshoot the problem. The storage system stores the core dump 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.

If the `coredump.savecore.warn` option is set to on, Data ONTAP prompts you for a confirmation when you delete unsaved core files.

Considerations for managing core dump files

A core dump produces a core dump file that contains the contents of system memory, including the system memory for the Performance Acceleration Module (PAM) or Flash Cache family of modules and the system memory for NVRAM. Several considerations exist for managing core dump files.

When a core dump file is created, it is stored in uncompressed format if sufficient space is available. If sufficient space is not available, it is stored in compressed format on a spare disk. If no spare disks are available, the system attempts to store the compressed core dump file across reserved sections of disks. You use the `coredump.dump.attempts` option to control how many attempts the system makes to create a core dump file. The default value is 2.

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.

You must not further compress a core dump file if you send it to technical support for analysis. The manual compression makes the file unrecognizable and can delay technical support's response time to your system issues. You can, however, use the `coredump segment config` and the `coredump segment` commands to segment the core dump file for easier handling.

Methods of segmenting core dump files

A core dump file can be very large, making it time consuming to upload to technical support when you need to. Segmenting the core dump file enables you to upload only the needed portion instead of the entire file.

You can segment a saved core dump file into a maximum of three core segments:

This core segment...	Contains system information from the memory of...
Primary core segment	Data ONTAP and the systemshell
Caching module core segment	Flash Cache family of modules
NVRAM core segment	NVRAM

Segmenting the core dump file enables you to upload a portion of the file as you need to. For instance, instead of uploading the entire core dump file to technical support for a core dump analysis, you can upload only the primary core segment of the file, and if necessary, upload the caching module core segment or NVRAM core segment later.

By using the `coredump segment config` commands, you can configure the automatic segmenting of the core dump file in the following ways:

- Specify whether to automatically segment a core dump file after it is saved
The default setting for automatic segmenting is system dependent.
- Specify whether to automatically delete the original core dump file after it is segmented
By default, automatic deletion of the original core dump file is disabled.
- Display the current configuration of the automatic segmenting of core dump files

By using the `coredump segment` commands, you can manually manage the segmenting of a core dump file in the following ways:

- Manually schedule a core segmenting job to segment a specified core dump file into core segments and specify whether the original core dump file is to be deleted after the core segmenting is complete
- Display information about core segments
- Delete a specified core segment or all segments
- Display the status of a core segmenting job
- Cancel a core segmenting job as specified by its job ID

Commands for managing core segmenting

You use the `coredump segment config` commands to manage the automatic segmenting of core dump files. You use the `coredump segment` commands to manage core segments.

If you want to...	Use this command...
Configure the automatic segmenting of core dump files, including: <ul style="list-style-type: none"> • Whether to automatically segment a core dump file after it is saved • Whether to automatically delete the original core dump file after it is segmented 	<code>coredump segment config modify</code>

If you want to...	Use this command...
Show the current configuration of automatic core segmenting	<code>coredump segment config show</code>
Manually start segmenting a specified core dump file into core segments and specify whether the original core dump file is to be deleted after the core segmenting is complete	<code>coredump segment start</code>
Display information about the core segments on the system, for example: <ul style="list-style-type: none"> • The core segment name • Total number of core segments for the full core • The time when the panic occurred that generated the core dump file 	<code>coredump segment show</code>
Delete a specified core segment from the system	<code>coredump segment delete</code>
Delete all core segments from the system	<code>coredump segment delete-all</code>
Displays the status of a core segmenting job, including the following: <ul style="list-style-type: none"> • Job ID • Name of the core dump file that is being segmented • Job status • Percent completed 	<code>coredump segment status</code>
Cancel a core segmenting job as specified by its job ID	<code>coredump segment stop</code>

For more information, see the man pages.

Automatic technical support notification upon system reboots

Your storage system sends email 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 email, you should contact technical support when your system creates a core file.

Understanding 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` file 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 about 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 following entities:

- The console
- A file
- A remote system

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

Related concepts

[How to access the default directories on the storage system](#) on page 79

[The `/etc/messages` file](#) on page 78

[How to access the default directories on the storage system](#) on page 79

Related tasks

[Accessing log files using HTTP or HTTPS](#) on page 82

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
<code>auth</code>	Messages from the authentication system, such as <code>login</code>
<code>cron</code>	Messages from the internal <code>cron</code> facility
<code>daemon</code>	Messages from storage system daemons, such as <code>rshd</code>
<code>kern</code>	Messages from the storage system kernel
<code>*</code>	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
<code>emerg</code>	Panic condition that causes a disruption of normal service
<code>alert</code>	Condition that you should correct immediately, such as a failed disk
<code>crit</code>	Critical conditions, such as disk errors
<code>err</code>	Errors, such as those caused by a bad configuration file
<code>warning</code>	Conditions that might become errors if not corrected
<code>notice</code>	Conditions that are not errors, but might require special handling
<code>info</code>	Information, such as the hourly uptime message
<code>debug</code>	Used for diagnostic purposes
<code>*</code>	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.	/etc/messages
Send messages to a host name preceded by an @ sign.	@adminhost
Send messages to the console.	/dev/console or *

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 142

Understanding 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 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 511 characters. An audit entry is truncated to 511 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 File Access and Protocols Management Guide for 7-Mode*.

For information about forwarding audit logs to a remote syslog log host, see the `na_auditlog(5)` man page.

Related concepts

[How to access the default directories on the storage system](#) on page 79

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.

Commands in the `/etc/rc` file

The `/etc/rc` file stores startup commands that your storage system executes at boot time to configure the system.

After you run the `setup` command or the Setup Wizard, the system automatically stores startup commands in the `/etc/rc` file.

The commands in the `/etc/rc` file configure the storage system to do the following:

- 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

The `/etc/rc` file must not contain the following types of commands:

- Commands that are executed by subsystems that are not yet available when the file is executed
- Commands that are interactive and would wait for input during the boot process

For example, you must not include the `iscsi` commands or the `wrfile` command in the `/etc/rc` file. Doing so prevents your storage system from booting successfully.

Running the `setup` command rewrites the `/etc/rc` file. It is recommended that you 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.

The following sample shows startup commands that are used in the `/etc/rc` file on the root volume:

```
#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 components of the sample `/etc/rc` file:

Description	Explanation
hostname toaster	Sets the storage system host name to “toaster”.
ifconfig e0 'hostname'-0 ifconfig e1 'hostname'-1 ifconfig f0 'hostname'-f0 ifconfig a5 'hostname'-a5	Sets the IP addresses for the storage system network interfaces with a default network mask. The arguments in single backquotes expand to “toaster” if you specify “toaster” as the host name during setup. The actual IP addresses are obtained from the <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.

Description	Explanation
<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 MyRouterBox 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 Network Management Guide for 7-Mode*.

Related concepts

[Managing core dump files](#) on page 139

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. Ensure that entries in the `/etc/rc` file are listed in the following order:

```
hostname system_name
ifgrp commands
vlan commands
ifconfig commands
vfiler commands
```

```
route commands
[any other commands]
```

4. Save the edited file.
5. 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 `/etc/rc` errors might cause the system to become inaccessible:

- An incorrect network address was specified using the `ifconfig` command.
The storage system is inaccessible because it is not on the network.
- Storage system directories were improperly exported 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.

When you back up a storage system configuration, the following files are backed up for the storage system and the default vFiler unit (vfiler0):

- System-specific configuration files, for example, `/etc/rc`
- System-specific registry options
- Volume configuration
- vfiler0-specific configuration, for example, `/etc/quotas`, `/etc/hosts`, `/etc/usermap.cfg`, `/etc/nsswitch.conf`, and `/etc/hosts.equiv`
- vfiler0-specific registry options, for example, NFS, CIFS, ndmpd, and NIS

If you have configured vFiler units, when you back up the configuration of a vFiler unit, the following files in the vFiler units are backed up:

- vFiler-specific configuration files, for example, `/etc/quotas`, `/etc/hosts`, `/etc/usermap.cfg`, `/etc/nsswitch.conf`, and `/etc/hosts.equiv`
- vFiler-specific registry options, for example, NFS, CIFS, ndmpd, and NIS

vFiler configuration is backed up or restored only for the vFiler unit on which the `config dump` or `config restore` command is run.

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.raidsizes`, 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.
3. If you use quotas for any volumes owned by a non-default vFiler unit (a vFiler unit other than `vfiler0`), ensure that the quotas are in the desired state (`on` or `off`) for those volumes.

The quotas state for volumes owned by a non-default vFiler is not restored when you restore a system configuration.

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.

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, you should enclose the line to be written within 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 154

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 153

Monitoring the storage system

You can use functionality such as event messages, health monitors, and AutoSupport to monitor the storage system.

Managing event messages

The Event Management System (EMS) collects and displays information about events that occur on your storage system. You can display the status of events. You can also display the event log and its contents.

Event messages appear on your system console or LCD, if your system has one, and are written to the system's event log. An event message consists of the following elements:

- Message name
- Severity level
 - Possible values include the following, listed in decreasing order of urgency:
 - EMERGENCY (the system is unusable)
 - ALERT (action must be taken immediately to prevent system failure)
 - CRITICAL
 - ERROR
 - WARNING
 - NOTICE (a normal but significant condition has occurred)
 - INFORMATIONAL
 - DEBUG
- Description

You can display the following information about events:

- Status of the events that have occurred on the system
 - You can also view system events through the system's RLM or SP
- The event log and its contents over a specified period of time

Displaying event information

You can display information about the status of the events that have occurred on the system.

Step

1. To display information about events that have occurred, use the `ems event status` command.

Example

The following example displays the output of the `ems event status` command.

```
Current time: 11May2011 20:37:31
Engine status: total 4220, drops 0, suppr (dup 0, timer 0, auto 0)
Event:Priority                Last Time
      Indications      Drops      DupSuppr      TimerSuppr      AutoSuppr
asup.general.reminder:INFO    5/11/2011 07:21:00
      2                0          0              0              0
callhome.management.log:INFO  5/11/2011 00:20:58
      1                0          0              0              0
callhome.nht.data:INFO        5/10/2011 08:00:00
      1                0          0              0              0
callhome.performance.data:INFO 5/11/2011 00:00:00
      1                0          0              0              0
...
```

Displaying event log information

You can view information about the event log and display its contents over a specified period of time.

Steps

1. To view information about the event log, use the `ems log status` command.

Example

The following example displays the output of the `ems log status` command.

```
EMS log data:
[LOG_default:
    enabled on, save 5, rotate weekly, size 9288893
    file /etc/log/ems, formal xml
    level debug
    indications 4230, drops 24
    last update: Wed, 11 May 2011 21:36:06 GMT
```

2. To display the contents of the event log, use the `ems log dump` command.

You specify the period of time by specifying the number of hours or days, as shown in the following examples.

```
ems log dump 4h
```

```
ems log dump 1d
```

Managing AutoSupport

AutoSupport is a mechanism that proactively monitors the health of your system and automatically sends email messages to NetApp technical support, your internal support organization, and a support partner.

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, but you should leave it enabled. Enabling AutoSupport can significantly help speed problem determination and resolution should a problem occur on your storage system. By default, the system collects AutoSupport information and stores it locally even if you disable AutoSupport.

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.

For more information about AutoSupport, see the NetApp Support Site.

Related information

The NetApp Support Site: support.netapp.com

When and where AutoSupport messages are sent

AutoSupport sends messages to different recipients, depending on the type of message. Learning when and where AutoSupport sends messages can help you understand messages that you receive through email or view on the My AutoSupport web site.

Note: Unless specified otherwise, settings in the following tables are options.

Event-triggered messages

When events occur on the storage system that require corrective action, AutoSupport automatically sends an event-triggered message.

When the message is sent	Where the message is sent
AutoSupport responds to a trigger event in the EMS	Addresses specified in <code>autosupport.to</code> and <code>autosupport.noteto</code> . (Only critical, service-affecting events are sent. The message sent to Note To is a shortened version.) Addresses specified in <code>autosupport.partner.to</code> Technical support, if <code>autosupport.support.enable</code> is on

Scheduled messages

AutoSupport automatically sends a number of messages on a regular schedule.

When the message is sent	Where the message is sent
Daily (log message)	Addresses specified in <code>autosupport.partner.to</code> Technical support, if <code>autosupport.support.enable</code> is on
Daily (performance message), if the <code>autosupport.performance_data.enable</code> option is on	Addresses specified in <code>autosupport.partner.to</code> Technical support, if <code>autosupport.support.enable</code> is on
Weekly (Sent Sunday between 12:00 a.m and 1:00 a.m.)	Addresses specified in <code>autosupport.partner.to</code> Technical support, if <code>autosupport.support.enable</code> is on

Manually triggered messages

You can manually initiate or resend an AutoSupport message.

When the message is sent	Where the message is sent
You manually initiate a message using the option <code>autosupport.doit text</code> command	Addresses specified in <code>autosupport.to</code> and <code>autosupport.partner.to</code>
You manually resend a past message using the <code>autosupport history retransmit</code> command	Only to the URI that you specify in the <code>-uri</code> parameter of the <code>autosupport history retransmit</code> command

Technical support triggered messages

Technical support can request messages from AutoSupport using the AutoSupport On Demand feature.

When the message is sent	Where the message is sent
When AutoSupport obtains delivery instructions to generate new AutoSupport messages	Addresses specified in <code>autosupport.partner.to</code> Technical support, if <code>autosupport.support.enable</code> is on and the transport protocol is HTTPS
When AutoSupport obtains delivery instructions to resend past AutoSupport messages	Technical support, if <code>autosupport.support.enable</code> is on and the transport protocol is HTTPS

Related concepts

[How AutoSupport On Demand obtains delivery instructions from technical support](#) on page 160

How event-triggered AutoSupport messages work

AutoSupport creates event-triggered AutoSupport messages when the EMS processes a trigger event. An event-triggered AutoSupport message alerts recipients of problems that require corrective action, and messages contain only information that is relevant to the problem. You can customize what content to include and who receives the messages.

AutoSupport uses the following process to create and send event-triggered AutoSupport messages:

1. When the EMS processes a trigger event, EMS sends AutoSupport a request.

Note: A trigger event is an EMS event with an AutoSupport destination and a name that begins with a `callhome.` prefix.

2. AutoSupport creates an event-triggered AutoSupport message.

AutoSupport collects basic and troubleshooting information from subsystems that are associated with the trigger to create a message that only includes information that is relevant to the trigger event.

A default set of subsystems are associated with each trigger. However, you can choose to associate additional subsystems with a trigger by using the `autosupport trigger modify` command.

3. AutoSupport sends the event-triggered AutoSupport message to the recipients defined by the `options autosupport.to`, `options autosupport.noteto`, `options autosupport.partner.to`, and `options autosupport.support.enable` commands. You can enable and disable delivery of AutoSupport messages for specific triggers by using the `autosupport trigger modify` command with the `-to` and `-noteto` parameters.

Example of data sent for a specific event

The `storage shelf PSU failed EMS` event triggers a message that contains basic data from the Mandatory, Log Files, Storage, RAID, HA, Platform, and Networking subsystems and troubleshooting data from the Mandatory, Log Files, and Storage subsystems.

You decide that you want to include data about NFS in any AutoSupport messages sent in response to a future `storage shelf PSU failed` event. You enter the following command to enable troubleshooting-level data for NFS for the `callhome.shlf.ps.fault` event:

```
system> autosupport trigger modify -autosupport-message
shlf.ps.fault -troubleshooting-additional nfs
```

Note: The `callhome.` prefix is omitted from the `storage shelf PSU failed` event when you use the `autosupport trigger` commands.

How AutoSupport On Demand obtains delivery instructions from technical support


AutoSupport On Demand periodically communicates with technical support to obtain delivery instructions for sending, resending, and declining AutoSupport messages. AutoSupport On Demand is enabled by default. It automatically communicates with technical support if AutoSupport is configured to send messages to technical support and the transport protocol is HTTPS.

The AutoSupport On Demand client, which runs on each node, periodically polls the AutoSupport On Demand service to obtain delivery instructions. The AutoSupport On Demand service resides in technical support. The client sends HTTPS requests to the same technical support location that AutoSupport messages are sent.

The AutoSupport On Demand client does not accept incoming connections.

Note: AutoSupport On Demand uses the "autosupport" user account to communicate with technical support. You should not delete this account.

The following illustration shows how AutoSupport On Demand sends HTTPS requests to technical support to obtain delivery instructions.



The delivery instructions can include requests for AutoSupport to do the following:

- Generate new AutoSupport messages.
Technical support might request new AutoSupport messages to help triage issues.
- Retransmit previously generated AutoSupport messages.
This request automatically happens if a message was not received due to a delivery failure.
- Disable delivery of AutoSupport messages for specific trigger events.
Technical support might disable delivery of data that is not used.

What data AutoSupport messages contain

AutoSupport messages contain information from subsystems. Learning what AutoSupport messages contain can help you interpret or respond to messages that you receive via email or view on the My AutoSupport web site.

Type of message	What type of data the message contains
Event-triggered	Files containing context-sensitive data about the specific subsystem where the event occurred
Daily	Log files
Performance	Performance data sampled during the previous 24 hours
Weekly	Configuration and status data
Triggered by the option <code>autosupport.doit text</code> command	<p>Depends on the text that you enter in the command:</p> <ul style="list-style-type: none"> • Including the word <code>test</code> sends a short message that triggers an automated response from the NetApp mail handler so that you can confirm that AutoSupport messages are being received. • Excluding the word <code>test</code> sends a set of data similar to the weekly message and includes troubleshooting data from each subsystem.
Triggered by AutoSupport On Demand	<p>AutoSupport On Demand can request new messages or past messages. The type of data included in those messages is as follows:</p> <p>New messages Depends on the type of AutoSupport collection, which can be <code>test</code>, <code>all</code>, or <code>performance</code>.</p> <p>Past messages Depends on the type of message that is resent.</p>

AutoSupport subsystems

Each subsystem provides basic and troubleshooting information that AutoSupport uses for its messages. Each subsystem is also associated with trigger events that allow AutoSupport to only collect information from subsystems that is relevant to the trigger event.

You can view information about subsystems and trigger events by using the `autosupport trigger show` command.

Files sent in event-triggered AutoSupport messages

Event-triggered AutoSupport messages only contain basic and troubleshooting information from subsystems that are associated with the event that caused AutoSupport to generate the message. The specific data helps you troubleshoot the problem.

AutoSupport uses the following criteria to control content in event-triggered AutoSupport messages:

- Which subsystems are included
Data is grouped into subsystems, including common subsystems, such as Log Files, and specific subsystems, such as RAID. Each event triggers a message that contains only the data from specific subsystems.
- The detail level of each included subsystem
Data for each included subsystem is provided at a basic or troubleshooting level.

You can view all possible events and determine which subsystems are included in messages about each event using the `autosupport trigger show` command with the `-instance` parameter.

In addition to the subsystems that are included by default for each event, you can add additional subsystems at either a basic or a troubleshooting level using the `autosupport trigger modify` command.

Log files sent in AutoSupport messages

AutoSupport messages can contain several key log files that enable technical support staff and your internal support organization to review recent system activity.

All types of AutoSupport messages include the following log files when the Log Files subsystem is enabled:

Log file	Amount of data included from the file
<ul style="list-style-type: none">• Log files from the <code>/mroot/etc/log/mlog/</code> directory• The MESSAGES log file	Only new lines added to the logs since the last AutoSupport message up to a specified maximum. This ensures that AutoSupport messages have unique, relevant—not overlapping—data. (Log files from partners are the exception; for partners, the maximum allowed data is included.)
<ul style="list-style-type: none">• Log files from the <code>/mroot/etc/log/shelflog/</code> directory• Log files from the <code>/mroot/etc/log/acp/</code> directory• Event Management System (EMS) log data	The most recent lines of data up to a specified maximum.

Files sent in weekly AutoSupport messages

Weekly AutoSupport messages contain additional configuration and status data that is useful to track changes in your system over time.

The following information is sent in weekly AutoSupport messages:

- Basic information about every subsystem
- Contents of selected `/mroot/etc` directory files
- Log files
- Output of commands that provide system information
- Additional information, including replicated database (RDB) information, service statistics, and more

Structure of AutoSupport messages sent via email

When an AutoSupport message is sent via email, the message has a standard subject, a brief body, and a large attachment in 7z file format that contains the data.

Note: If AutoSupport is configured to hide private data, certain information, such as the hostname, is omitted or masked in the header, subject, body, and attachments.

Subject

The subject line of messages sent by the AutoSupport mechanism contains a text string that identifies the reason for the notification. The format of the subject line is as follows:

System Notification from *System_Name (Message) Severity*, where:

- "System" is replaced with "HA Group" if the storage system is configured for high availability
- *System_Name* is either the hostname or the system ID, depending on the AutoSupport configuration

Body

The body of the AutoSupport message contains the following information:

- Date and timestamp of the message
- Version of Data ONTAP
- System ID, serial number, and hostname
- AutoSupport sequence number
- SNMP contact name and location, if specified
- System ID and hostname of the HA partner, if the storage system is configured for high availability

Attached files

The key information in an AutoSupport message is contained in files that are compressed together into a 7z file called `body.7z` and attached to the message.

The files contained in the attachment are specific to the type of AutoSupport message.

AutoSupport severity types

AutoSupport messages have severity types that help you understand the purpose of each message—for example, to draw immediate attention to a critical problem, or only to provide information.

Messages have one of the following severities:

- Critical: critical conditions
- Error: error conditions
- Warning: warning conditions
- Notice: normal but significant condition
- Info: informational message
- Debug: debug-level messages

If your internal support organization receives AutoSupport messages via email, the severity appears in the subject line of the email message.

AutoSupport transport protocols

AutoSupport supports HTTPS, HTTP, and SMTP as the transport protocols for delivering AutoSupport messages to NetApp technical support. All of these protocols run on IPv4 or IPv6 based on the address family the name resolves to. If you enable AutoSupport messages to your internal support organization, those messages are sent by SMTP.

Protocol availability varies with the destination of the AutoSupport messages:

- If you enable AutoSupport to send messages to NetApp technical support, you can use any of the following transport protocols:

Protocol and port	Description
HTTPS on port 443	<p>This is the default protocol. You should use this whenever possible.</p> <p>The certificate from the remote server is validated against the root certificate, unless you disable validation.</p> <p>The delivery uses an HTTP PUT request. With PUT, if the request fails during transmission, the request restarts where it left off. If the server receiving the request does not support PUT, the delivery uses an HTTP POST request.</p>

Protocol and port	Description
HTTP on port 80	This protocol is preferred over SMTP. The delivery uses an HTTP PUT request. With PUT, if the request fails during transmission, the request restarts where it left off. If the server receiving the request does not support PUT, the delivery uses an HTTP POST request.
SMTP on port 25	You should use this protocol only if the network connection does not allow HTTPS or HTTP, because SMTP can introduce limitations on message length and line length.

- If you configure AutoSupport with specific email addresses for your internal support organization, those messages are always sent by SMTP.

For example, if you use the recommended protocol to send messages to NetApp technical support and you also want to send messages to your internal support organization, your messages would be transported via both HTTPS and SMTP, respectively.

AutoSupport limits the maximum file size for each protocol. The default setting for HTTP and HTTPS transfers is 10 MB. The default setting for SMTP transfers is 5 MB. If the size of the AutoSupport message exceeds the configured limit, AutoSupport delivers as much of the message as possible. You can edit the maximum size by modifying AutoSupport configuration. See the `options` man page for more information.

The protocols require the following additional configuration:

- If you use HTTP or HTTPS to send AutoSupport messages to NetApp technical support and you have a proxy, you must identify the URL for that proxy.
If the proxy uses a port other than the default port, which is 3128, you can specify the port for that proxy. You can also specify a username and password for proxy authentication.
- If you use SMTP to send AutoSupport messages either to your internal support organization or to NetApp technical support, you must have an external mail server.

The storage system does not function as a mail server—it requires an external mail server at your site to send mail. The mail server must be a host that listens on the SMTP port (25), and it must be configured to send and receive 8-bit Multipurpose Internet Mail Extensions (MIME) encoding. Example mail hosts include a UNIX host running an SMTP server such as the `sendmail` program and a Windows NT server running the Microsoft Exchange server. You can have one or more mail hosts.

No matter what transport protocol you use, you can use IPv4 or IPv6 addresses based on the address family to which the name resolves.

Setting up AutoSupport

You can control whether and how AutoSupport information is sent to NetApp technical support and your internal support organization, and then test that the configuration is correct.

About this task

For more information about the following commands, see the man pages.

Steps

1. Ensure AutoSupport is enabled by setting the `autosupport.enable` option to `on`.
2. If you want technical support to receive AutoSupport messages, set the following options:
 - a) Set `autosupport.support.enable` to `on`.
 - b) Select a transport protocol for messages to NetApp technical support by setting `autosupport.support.transport` to `smtp`, `http`, or `https`.
 - c) If you chose HTTP or HTTPS as the transport protocol and you use a proxy, set `autosupport.proxy.url` to the URL of your proxy.
3. If you want your internal support organization or a support partner to receive AutoSupport messages, perform the following actions:
 - a) Identify the recipients in your organization by setting the following options:

Set this option	To this
<code>autosupport.to</code>	Up to five comma-separated individual email addresses or distribution lists in your internal support organization that will receive key AutoSupport messages
<code>autosupport.noteto</code>	Up to five comma-separated individual email addresses or distribution lists in your internal support organization that will receive a shortened version of key AutoSupport messages designed for cell phones and other mobile devices
<code>autosupport.partner.to</code>	Up to five comma-separated individual email addresses or distribution lists in your support partner organization that will receive all AutoSupport messages

- b) Check that addresses are correctly configured by listing the destinations using the `autosupport destinations show` command.
4. If you are sending messages to your internal support organization or you chose SMTP transport for messages to technical support, configure SMTP by setting the following options:

- Set `autosupport.mailhost` to one or more mail hosts, separated by commas. You can set a maximum of five.
 - Set `autosupport.from` to the email address that sends the AutoSupport message.
 - Set `autosupport.max_smtp_size` to the email size limit of your SMTP server.
5. If you want AutoSupport to specify a fully qualified domain name when it sends connection requests to your SMTP mail server, configure DNS.

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

6. Optional: Change the following settings:

If you want to do this...	Set the following options...
Hide private data by removing, masking, or encoding sensitive data in the messages	Set <code>autosupport.content</code> to <code>minimal</code> . Note: If you change from <code>complete</code> to <code>minimal</code> , all AutoSupport history and all associated files are deleted.
Stop sending performance data in periodic AutoSupport messages	Set <code>autosupport.performance_data.enable</code> to <code>disable</code> .

7. Check the overall configuration using the `options autosupport` command.
8. Test that AutoSupport messages are being sent and received:
- a) Use the `options autosupport.doit test` command.
 - b) Confirm that NetApp is receiving your AutoSupport messages by checking the email address that technical support has on file for the system owner, who should have received an automated response from the NetApp mail handler.
 - c) Optional: Confirm that the AutoSupport message is being sent to your internal support organization or to your support partner by checking the email of any address that you configured for the `autosupport.to`, `autosupport.noteto`, or `autosupport.partner` options.

Related tasks

[Troubleshooting AutoSupport when messages are not received](#) on page 177

Getting AutoSupport message descriptions

The descriptions of the AutoSupport messages that you receive are available through the online AutoSupport Message Matrices page.

Steps

1. Go to the AutoSupport Message Matrices page: support.netapp.com/NOW/knowledge/docs/olio/autosupport/matrices/
2. On the AutoSupport Message Matrices page under Select a Release, select your version of Data ONTAP and click **View Matrix**.

The Syslog Translator page appears with all AutoSupport message descriptions listed alphabetically by subject line.

Commands for managing AutoSupport

You use the `autosupport` and `options autosupport` commands to change or view AutoSupport configuration, display information about past AutoSupport messages, and send or resend an AutoSupport message.

Configure AutoSupport

If you want to...	Use this command...
Control whether any AutoSupport messages are sent	<code>options autosupport.enable</code>
Control whether AutoSupport messages are sent to technical support	<code>options autosupport.support.enable</code>
Set up AutoSupport or modify the configuration of AutoSupport	<code>options autosupport</code>
Enable and disable AutoSupport messages to your internal support organization for individual trigger events, and specify additional subsystem reports to include in messages sent in response to individual trigger events	<code>autosupport trigger modify</code>

Display information about the configuration of AutoSupport

If you want to...	Use this command...
Display the AutoSupport configuration	<code>options autosupport</code>
View a summary of all addresses and URLs that receive AutoSupport messages	<code>autosupport destinations show</code>

If you want to...	Use this command...
Display which AutoSupport messages are sent to your internal support organization for individual trigger events	<code>autosupport trigger show</code>

Display information about past AutoSupport messages

If you want to...	Use this command...
Display information about one or more of the 50 most recent AutoSupport messages	<code>autosupport history show</code>
View the information in the AutoSupport messages including the name and size of each file collected for the message along with any errors	<code>autosupport manifest show</code>

Send or resend AutoSupport messages

If you want to...	Use this command...
Retransmit a locally stored AutoSupport message, identified by its AutoSupport sequence number Note: If you retransmit an AutoSupport message, and if support already received that message, the support system will not create a duplicate case. If, on the other hand, support did not receive that message, then the AutoSupport system will analyze the message and create a case, if necessary.	<code>autosupport history retransmit</code>
Generate and send an AutoSupport message—for example, for testing purposes	<code>options autosupport.doit</code>

For more information, see the man pages.

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)` man page.

AutoSupport option	Description
<code>autosupport.content</code> [complete minimal]	<p>Indicates whether private data is omitted from AutoSupport messages.</p> <p>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>Setting this option to <code>minimal</code> removes, encodes, or masks sensitive data from AutoSupport attachments and headers. The affected data might include IP addresses, MAC addresses, URIs, DNS names, email addresses, port numbers, filer names, aggregate names, volume names, junction paths, policy names, user IDs, group IDs, LUNs, and qtree names.</p> <p>Note: If you change from <code>complete</code> to <code>minimal</code>, all AutoSupport history and all associated files are deleted.</p>
<code>autosupport.doit</code> [message]	<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>

AutoSupport option	Description
<code>autosupport.from sender</code>	<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_collection</code>	<p>Enables and disables local storage of AutoSupport files when sending of AutoSupport messages is disabled. The default setting is <code>true</code>, which causes the node to store AutoSupport files locally even if AutoSupport is disabled.</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 between entries. The default is "mailhost."</p> <p>The specified mail hosts will be used to send AutoSupport messages to all email address specified in other parameter. The specified mail hosts are used to send AutoSupport messages to NetApp technical support if <code>autosupport.support.transport</code> is <code>smtp</code>. Both IPv6 and IPv4 addresses are accepted.</p>
<code>autosupport.max_http_size</code> <i>integer[KB MB GB TB PB]</i>	<p>Specifies the maximum file size for HTTP and HTTPS transfers of AutoSupport messages to NetApp technical support if <code>autosupport.support.transport</code> is <code>http</code> or <code>https</code>. The default is 10MB.</p>
<code>autosupport.max_smtp_size</code> <i>integer[KB MB GB TB PB]</i>	<p>Specifies the maximum email message size for SMTP (email) transfers of AutoSupport messages. This option applies to all messages sent to <code>autosupport.to</code>, <code>autosupport.noteto</code>, and <code>autosupport.partner.to</code>. It also applies to messages sent to NetApp technical support if <code>autosupport.support.transport</code> is <code>smtp</code>. You should set this value to the email size limit of your SMTP server. The default is 5MB.</p>

AutoSupport option	Description
<code>autosupport.minimal.subject.id</code> [<i>hostname</i> <i>systemid</i>]	Defines how the system is identified in the AutoSupport message title if <code>autosupport.content</code> is minimal. The default is <code>systemid</code> .
<code>autosupport.nht_data.enable</code>	This option is no longer used. AutoSupport no longer sends NHT disk drive data in AutoSupport messages.
<code>autosupport.noteto address1[, ..., address5]</code>	<p>Defines the list of recipients for the AutoSupport short note email. The short note email 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 email addresses are allowed. Enter the addresses as a comma-separated list with no spaces between entries. The default is an empty list, which disables short note emails.</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.ondemand.polling_interval</code>	Defines the rate in which the node polls the AutoSupport On Demand service. This option is not editable at the admin privilege level. It displays for informational purposes only.
<code>autosupport.ondemand.remotediag.state</code>	Defines whether the AutoSupport On Demand Remote Diagnostics feature is enabled or disabled on the node. The default is <code>on</code> . This option is not editable at the admin privilege level. It displays for informational purposes only.
<code>autosupport.ondemand.server_url</code>	Defines the AutoSupport On Demand service URL that the node communicates with. This option is not editable at the admin privilege level. It displays for informational purposes only.
<code>autosupport.ondemand.state</code>	Defines whether the AutoSupport On Demand feature is enabled or disabled on the node. The default is <code>on</code> . This option is not editable at the admin privilege level. It displays for informational purposes only.

AutoSupport option	Description
<code>autosupport.partner.to address1[, ..., address5]</code>	<p>Defines the list of recipients who will receive all AutoSupport email notifications regardless of the severity level.</p> <p>Up to five email addresses are allowed. Enter the addresses as a comma-separated list with no spaces between entries. 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.payload.format</code>	<p>Specifies the file format of the compressed file that contains AutoSupport data. Use <code>7z</code> to specify 7-Zip archive format. Use <code>tgz</code> to specify GNU zipped tar file. The default is <code>7z</code>.</p>
<code>autosupport.performance_data.doit any_string</code>	<p>Triggers a performance snapshot AutoSupport message when any string is added.</p>
<code>autosupport.performance_data.enable</code>	<p>Enables sending messages about performance data to technical support and addresses specified in <code>autosupport.partner.to</code>. This option should always be set to <code>on</code>. The default is <code>on</code>.</p>
<code>autosupport.periodic.tx_window time</code>	<p>Specifies the randomized delay window for periodic AutoSupport messages. Values can range from 0 to 240 minutes (4 hours). The default is 60 (1 hour). Setting the value to 0 disables the randomized delay, which is intended to prevent bottlenecks.</p>
<code>autosupport.retry.count #retries</code>	<p>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.</p>
<code>autosupport.retry.interval interval</code>	<p>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).</p>

AutoSupport option	Description
<code>autosupport.support.enable [on off]</code>	<p>Enables and disables sending of all AutoSupport messages to technical support. The default is on.</p>
<code>autosupport.support.proxy [user:pass@]proxyhost.com[:port][/]]</code>	<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>Both IPv6 and IPv4 addresses are accepted.</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>
<code>autosupport.support.put_url URL</code>	<p>Indicates where AutoSupport messages for NetApp technical support are sent if <code>autosupport.support.transport</code> is <code>http</code> or <code>https</code>. Each message sent via HTTP or HTTPS is sent as an HTTP PUT request to this URL. If the server receiving the message does not support PUT requests, the message is sent via HTTP POST to the URL indicated by <code>autosupport.support.url</code>. This option is read-only and is shown for informational purposes only.</p>

AutoSupport option	Description
<code>autosupport.support.reminder [on off]</code>	Enables or disables a reminder message that appears when <code>autosupport.support.enable</code> is set to <code>off</code> . The default is <code>on</code> .
<code>autosupport.support.to</code>	Indicates where AutoSupport messages for NetApp technical support are sent if <code>autosupport.support.transport</code> is <code>smtp</code> . This option is read-only and is shown for informational purposes only.
<code>autosupport.support.transport [http https smtp]</code>	Defines the type of delivery for AutoSupport messages that are sent to NetApp technical support. The default is <code>https</code> .
<code>autosupport.support.url URL</code>	Indicates where AutoSupport messages for NetApp technical support are sent if <code>autosupport.support.transport</code> is <code>http</code> or <code>https</code> . Each message sent via HTTP or HTTPS is sent as an HTTP PUT request to the URL indicated by <code>autosupport.support.put_url</code> . If the server receiving the message does not support PUT requests, the message is sent via HTTP POST to the URL indicated by this option. This option is read-only and is shown for informational purposes only.
<code>autosupport.throttle [on off]</code>	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> .

AutoSupport option	Description
<code>autosupport.to address1[, ..., address5]</code>	<p>Defines the list of recipients who receive key AutoSupport messages, as defined in factory-default settings. You can use the <code>autosupport trigger show</code> command to display AutoSupport trigger configuration and the <code>autosupport trigger modify</code> command to modify the trigger configuration.</p> <p>Up to five email addresses are allowed, or the list can be left empty. Enter the addresses as a comma-separated list with no spaces between entries. 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>
<code>autosupport.validate_digital_certificate [on off]</code>	<p>Determines whether the system validates remote digital certificates that it receives. Applies only when <code>autosupport.support.transport</code> is set to HTTPS.</p>

What the AutoSupport manifest is

The AutoSupport manifest provides a detailed view of the files collected for each event-triggered AutoSupport message. The AutoSupport manifest also includes information about collection errors when AutoSupport cannot collect the files it needs.

The AutoSupport manifest includes the following information:

- Sequence number of the event-triggered AutoSupport message
- What files AutoSupport included in the event-triggered AutoSupport message
- Size of each file in bytes
- Status of the AutoSupport manifest collection
- Error description if AutoSupport failed to collect one or more files

You can view the AutoSupport manifest by using the `autosupport manifest show` command.

Troubleshooting AutoSupport

If you do not receive AutoSupport messages, you can check a number of settings to resolve the problem.

Troubleshooting AutoSupport when messages are not received

If the system does not send the AutoSupport message, you can determine whether that is because AutoSupport cannot generate the message or cannot deliver the message.

Steps

1. Check delivery status of the messages by using the `autosupport history show` command.
2. Read the status.

This status	Means
initializing	The collection process is starting. If this state is temporary, all is well. However, if this state persists, there is an issue.
collection-failed	AutoSupport cannot create the AutoSupport content in the spool directory. You can view what AutoSupport is trying to collect by entering the <code>autosupport history show -detail</code> command.
collection-in-progress	AutoSupport is collecting AutoSupport content. You can view what AutoSupport is collecting. Obtain the sequence number by using the <code>autosupport history show</code> command, and then display the manifest details for the sequence number by entering the <code>autosupport manifest show</code> command with the <code>-seq-num</code> parameter.
queued	AutoSupport messages are queued for delivery, but not yet delivered.
transmitting	AutoSupport is currently delivering messages.
sent-successful	AutoSupport successfully delivered the message. You can find out where AutoSupport delivered the message by entering the <code>autosupport history show -delivery</code> command.
ignore	AutoSupport has no destinations for the message. You can view the delivery details by entering the <code>autosupport history show -delivery</code> command.
re-queued	AutoSupport tried to deliver messages, but the attempt failed. As a result, AutoSupport placed the messages back in the delivery queue for another attempt. You can view the error by entering the <code>autosupport history show</code> command.
transmission-failed	AutoSupport failed to deliver the message the specified number of times and stopped trying to deliver the message. You can view the error by entering the <code>autosupport history show</code> command.
ondemand-ignore	The AutoSupport message was processed successfully, but the AutoSupport On Demand service chose to ignore it.

3. Perform one of the following actions:

For this status	Do this
initializing or collection-failed	Contact technical support because AutoSupport cannot generate the message.
ignore, re-queued, or transmission failed	Check that destinations are correctly configured for SMTP, HTTP, or HTTPS because AutoSupport cannot deliver the message.

Related tasks

[Troubleshooting AutoSupport over SMTP](#) on page 179

[Troubleshooting AutoSupport over HTTP or HTTPS](#) on page 178

Troubleshooting AutoSupport over HTTP or HTTPS

If the system does not send the expected AutoSupport message and you are using HTTP or HTTPS, you can check a number of settings to resolve the problem.

Before you begin

You determined that AutoSupport can generate the message, but not deliver the message over HTTP or HTTPS.

Steps

1. At the storage system's CLI, ensure that DNS is enabled and configured correctly by entering the following command:

```
dns info
```

2. Read the error for the AutoSupport message by using the `autosupport history show` command with the `-seq-num` and `-destination` parameters.
3. At the storage system's CLI, ensure that the system is routing out to the Internet successfully by entering the following command:

```
traceroute -p port support.netapp.com
```

The default `port` is 80 for HTTP and 443 for HTTPS.

Note: If AutoSupport is configured to use a proxy, use the `traceroute -p` command to test the path to the proxy.

4. Use the `rdfile` command to read the `/etc/log/mlog/notifyd.log` file.

Related tasks

[Troubleshooting AutoSupport when messages are not received](#) on page 177

Troubleshooting AutoSupport over SMTP

If the system does not send the AutoSupport message and you are using SMTP, you can check a number of settings to resolve the problem.

Before you begin

You determined that AutoSupport can generate the message, but not deliver the message over SMTP.

Steps

1. At the clustershell CLI, ensure that DNS for the cluster is enabled and configured correctly by entering the following command:

```
dns info
```

2. At the clustershell CLI, check that the mail host specified in the configuration is a host that the storage system can talk to by entering the following command:

```
ping mailhost
```

mailhost is the name or IP address of your mail host.

3. 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.
```

4. At the CLI for the storage system, ensure that the system is reaching the mail host successfully by entering the following command:

```
traceroute -p mailhost
```

mailhost is the name or IP address of your mail host.

5. From some other host, telnet to the SMTP port by entering the following command:

```
telnet mailhost 25
```

A message similar to the following text is displayed:

```
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. Use the `rdfile` command to read the `/etc/log/mlog/notifyd.log` file.

Related tasks

[Troubleshooting AutoSupport when messages are not received](#) on page 177

Troubleshooting EMS events about rejected or failed SMTP attempts

If the system attempted to send an AutoSupport email, but the attempt resulted in an EMS event about a rejected or failed SMTP or an unknown user, you can check the relaying configuration for the mail host to determine whether relaying is denied or incorrectly configured.

About this task

The EMS identifiers for this event are `asup.smtp.fail` and `asup.smtp.reject`. You can use the EMS identifiers to view a description of the messages in the Syslog Translator on the NetApp Support Site.

Steps

1. From a Windows, UNIX, or Linux host, telnet to port 25 of the mail host by entering the following command:

```
telnet mailhost 25
```

2. Test whether relaying is denied on the mail host.

- a) Enter the following commands:

```
HELO DOMAIN NAME
```

```
MAIL FROM: your_email_address
```

```
RCPT TO: autosupport@netapp.com
```

- b) If you receive a message similar to `relaying denied`, contact the mail host vendor because relaying is denied. Otherwise, continue to the next step.

3. Test whether relaying is incorrectly configured on the mail host.

- a) Enter the following commands:

```
DATA
```

```
SUBJECT: TESTING
```

```
THIS IS A TEST
```

```
.
```

Note: Ensure that you enter the last period (.) on a line by itself. The period indicates to the mail host that the message is complete.

- b) If you receive a message similar to `unknown user` or `unknown mailbox`, contact the mail host vendor because relaying is incorrectly configured.

Monitoring the health of your system

Health monitors proactively monitor certain critical conditions and raise alerts if they detect a fault or risk. If there are active alerts, the system health status reports a degraded status for the cluster. The alerts include the information that you need to respond to degraded system health.

If the status is degraded, you can view details about the problem, including the probable cause and recommended recovery actions. After you resolve the problem, the system health status automatically returns to OK.

The system health status reflects multiple separate health monitors. A degraded status in an individual health monitor causes a degraded status for the overall system health.

How health monitoring works

Individual health monitors have a set of health policies that trigger alerts when certain conditions or state changes occur. Understanding how health monitoring works can help you respond to problems and control future alerts.

Health monitoring consists of the following components:

- Individual health monitors for specific subsystems, each of which has its own health status
For example, the Storage subsystem has a node connectivity health monitor.
- An overall system health monitor that consolidates the health status of the individual health monitors
A degraded status in any single subsystem results in a degraded status for the entire system. If no subsystems have alerts, the overall system status is OK.

Each health monitor is made up of the following key elements:

- Alerts that the health monitor can potentially raise
Each alert has a definition, which includes details such as the severity of the alert and its probable cause.
- Health policies that identify when each alert is triggered
Each health policy has a rule expression, which is the exact condition or change that triggers the alert.

A health monitor continuously monitors and validates the resources in its subsystem for condition or state changes. When a condition or state change matches a rule expression in a health policy, the health monitor raises an alert. An alert causes the subsystem's health status and the overall system health status to become degraded.

How you can respond to system health alerts

When a system health alert occurs, you can acknowledge it, learn more about it, repair the underlying condition, and prevent it from occurring again.

When a health monitor raises an alert, you can respond in any of the following ways:

- Get information about the alert, which includes the affected resource, alert severity, probable cause, possible effect, and corrective actions.
- Get detailed information about the alert, such as the time when the alert was raised and whether anyone else has acknowledged the alert already.
- Get health-related information about the state of the affected resource or subsystem, such as a specific shelf or disk.
- Acknowledge the alert to indicate that someone is working on the problem, and identify yourself as the "Acknowledger."
- Resolve the problem by taking the corrective actions provided in the alert, such as fixing cabling to resolve a connectivity problem.
- Delete the alert, if the system did not automatically clear it.
- Suppress an alert to prevent it from affecting the health status of a subsystem.

Suppressing is useful when you understand a problem. After you suppress an alert, it can still occur, but the subsystem health displays as "ok-with-suppressed" when the suppressed alert occurs.

How you can control when system health alerts occur

You can control which alerts a health monitor generates by enabling and disabling the system health policies that define when alerts are triggered. This enables you to customize the health monitoring system for your particular context.

You can learn the name of a policy either by displaying detailed information about a generated alert or by displaying policy definitions for a specific health monitor or alert ID.

Disabling health policies is different from suppressing alerts. When you suppress an alert, it doesn't affect the subsystem's health status, but the alert can still occur.

If you disable a policy, the condition or state that is defined in its policy rule expression no longer triggers an alert.

Example of an alert that you want to disable

For example, suppose an alert occurs that is not useful to you. You use the `system health alert show -instance` command to obtain the Policy ID for the alert. You use the policy ID in the `system health policy definition show` command to view information about the policy. After reviewing the rule expression and other information about the policy, you decide to disable the policy. You use the `system health policy definition modify` command to disable the policy.

How health alerts trigger AutoSupport messages and events

System health alerts trigger AutoSupport messages and events in the Event Management System (EMS), making it possible to monitor the health of the system using AutoSupport messages and the EMS in addition to using the health monitoring system directly.

Your system sends an AutoSupport message within five minutes of an alert. The AutoSupport message includes all alerts generated since the last AutoSupport message, except for alerts that duplicate an alert for the same resource and probable cause within the last week.

Some alerts do not trigger AutoSupport messages. An alert does not trigger an AutoSupport message if its health policy disables the sending of AutoSupport messages. For example, a health policy might disable AutoSupport messages by default because AutoSupport already generates a message when the problem occurs. You can configure policies to not trigger AutoSupport messages by using the `system health policy definition modify` command.

You can view a list of all of the alert-triggered AutoSupport messages sent in the last week using the `system health autosupport trigger history show` command.

Alerts also trigger the generation of events to the EMS. An event is generated each time an alert is created and each time an alert is cleared.

What health monitors are available

There are several health monitors that monitor different parts of a system.

Health monitor name (identifier)	Subsystem name (identifier)	Purpose
Node connectivity (node-connect)	CIFS non-disruptive operations (CIFS-NDO)	Monitors SMB connections to ensure non-disruptive operations to Hyper-V applications.
	Storage (SAS-connect)	Monitors shelves, disks, and adapters at the node level to ensure that they have appropriate paths and connections.
System	n/a	Aggregates information from other health monitors.

Responding to degraded system health

When your system's health status is degraded, you can show alerts, read about the probable cause and corrective actions, show information about the degraded subsystem, and resolve the problem.

About this task

You can discover that an alert was generated by viewing an AutoSupport message, an EMS event, or by using the `system health` commands.

Steps

1. Use the `system health alert show` command to view the alerts that are compromising the system's health.
2. Read the alert's probable cause, possible effect, and corrective actions to determine if you can resolve the problem or if you need more information.
3. If you need more information, take any of the following actions:
 - Use the `system health alert show -instance` command to view additional information available for the alert.
 - Use the specific commands in the `system health` command directory for the affected subsystem to investigate the problem.

Example

For example, if a disk has a problem, use the `system health node-connectivity disk` command to get more information about the disk.

4. Optional: Use the `system health alert modify` command with the `-acknowledge` parameter to indicate that you are working on a specific alert.
5. Take corrective action to resolve the problem as described by the Corrective Actions field in the alert.

The Corrective Actions might include rebooting the system.

When the problem is resolved, the alert is automatically cleared. If the subsystem has no other alerts, the health of the subsystem changes to OK. If the health of all subsystems is OK, the overall system health status changes to OK.

6. Use the `system health status show` command to confirm that the system health status is OK.

If the system health status is not OK, repeat this procedure.

Example of responding to degraded system health

By reviewing a specific example of degraded system health caused by a shelf that lacks two paths to a node, you can see what the CLI displays when you respond to an alert.

After starting Data ONTAP, you check the system health and you discover that the status is degraded.

```
system>system health status show
Status
-----
degraded
```

You show alerts to find out where the problem is, and see that shelf 2 does not have two paths to node1.


```

system>system health alert show
      Node: node1
      Resource: Shelf ID 2
      Severity: Major
      Probable Cause: Disk shelf 2 does not have two paths to controller
                      node1.
      Possible Effect: Access to disk shelf 2 via controller node1 will be
                      lost with a single hardware component failure (e.g.
                      cable, HBA, or IOM failure).
      Corrective Actions: 1. Halt controller node1 and all controllers attached to disk shelf 2.
                        2. Connect disk shelf 2 to controller node1 via two paths following the
                           rules in the Universal SAS and ACP Cabling Guide.
                        3. Reboot the halted controllers.
                        4. Contact support personnel if the alert persists.

```

You display details about the alert to get more information, including the alert ID.

```

system>system health alert show -monitor node-connect -alert-id DualPathToDiskShelf_Alert -
instance
      Node: node1
      Monitor: node-connect
      Alert ID: DualPathToDiskShelf_Alert
      Alerting Resource: 50:05:0c:c1:02:00:0f:02
      Subsystem: SAS-connect
      Indication Time: Mon Mar 21 10:26:38 2011
      Perceived Severity: Major
      Probable Cause: Connection_establishment_error
      Description: Disk shelf 2 does not have two paths to controller node1.
      Corrective Actions: 1. Halt controller node1 and all controllers attached to disk shelf 2.
                        2. Connect disk shelf 2 to controller node1 via two paths following
                           the rules in the Universal SAS and ACP Cabling Guide.
                        3. Reboot the halted controllers.
                        4. Contact support personnel if the alert persists.
      Possible Effect: Access to disk shelf 2 via controller node1 will be lost with a single
                      hardware component failure (e.g. cable, HBA, or IOM failure).
      Acknowledge: false
      Suppress: false
      Policy: DualPathToDiskShelf_Policy
      Acknowledger: -
      Suppressor: -
      Additional Information: Shelf uuid: 50:05:0c:c1:02:00:0f:02
                           Shelf id: 2
                           Shelf Name: 4d.shelf2
                           Number of Paths: 1
                           Number of Disks: 6
                           Adapter connected to IOMA:
                           Adapter connected to IOMB: 4d
      Alerting Resource Name: Shelf ID 2

```

You acknowledge the alert to indicate that you are working on it.

```

system>system health alert modify -node node1 -alert-id DualPathToDiskShelf_Alert -
acknowledge true

```

You fix the cabling between shelf 2 and node1, and reboot the system. Then you check system health again, and see that the status is OK.

```

system>system health status show
      Status

```

OK

Commands for monitoring the health of your system

You can use the `system health` commands to display information about the health of system resources, to respond to alerts, to configure future alerts, and to display information about how health monitoring is configured.

Displaying health status

If you want to...	Use this command...
Display the health status of the system, which reflects the overall status of individual health monitors	<code>system health status show</code>
Display the health status of subsystems for which health monitoring is available	<code>system health subsystem show</code>

Displaying the status of node connectivity

If you want to...	Use this command...
Display the status of shelves from the node-level view, along with other information, such as the owner node, shelf name, and how many disks and paths the shelf has	<code>system health node-connectivity shelf show</code> Note: Use the <code>-instance</code> parameter to display detailed information about each shelf.
Display the status of disks, along with other information, such as the owner node, disk name and bay number, and the number of paths to the disk	<code>system health node-connectivity disk show</code> Note: Use the <code>-instance</code> parameter to display detailed information about each disk.
Display the status of adapters, along with other information, such as the owner node, whether they are used and enabled, and the number of shelves attached	<code>system health node-connectivity adapter show</code> Note: Use the <code>-instance</code> parameter to display detailed information about each adapter.

Responding to generated alerts

If you want to...	Use this command...
Display information about generated alerts, such as the resource and node where the alert was triggered, and the alert's severity and probable cause.	<pre>system health alert show</pre> <p>Note: Use the <code>-instance</code> parameter to display detailed information about each generated alert. Use other parameters to filter the list of alerts—for example, by node, resource, severity, and so on.</p>
Indicate that someone is working on an alert	<pre>system health alert modify</pre> <p>with the <code>-acknowledge</code> parameter</p>
Suppress a subsequent alert so that it does not affect the health status of a subsystem	<pre>system health alert modify</pre> <p>with the <code>-suppress</code> parameter</p>
Delete an alert that was not automatically cleared	<pre>system health alert delete</pre>
Display information about the AutoSupport messages that alerts triggered within the last week—for example, to determine if an alert triggered an AutoSupport message	<pre>system health autosupport trigger history show</pre>

Configuring future alerts

If you want to...	Use this command...
Enable or disable the policy that controls whether a specific resource state raises a specific alert	<pre>system health policy definition modify</pre>

Displaying information about how health monitoring is configured

If you want to...	Use this command...
Display information about health monitors, such as their nodes, names, subsystems, and status	<pre>system health config show</pre> <p>Note: Use the <code>-instance</code> parameter to display detailed information about each health monitor.</p>

If you want to...	Use this command...
Display information about the alerts that a health monitor can potentially generate	<code>system health alert definition show</code> Note: Use the <code>-instance</code> parameter to display detailed information about each alert definition.
Display information about health monitor policies, which determine when alerts are raised	<code>system health policy definition show</code> Note: Use the <code>-instance</code> parameter to display detailed information about each policy. Use other parameters to filter the list of alerts—for example, by policy status (enabled or not), health monitor, alert, and so on.

For more information, see the man pages for the commands.

Managing a storage system remotely

You can manage a storage system remotely by using a remote management device, which can be the SP or the RLM, depending on the platform model. The device stays operational regardless of the operating state of the system. You can also download the RSA as an upgrade to the remote management device.

The RLM is included in the 31xx, 6040, and 6080 platforms.

The SP is included in all other platform models.

Additionally, you can download the Remote Support Agent (RSA), a firmware upgrade to the SP and the RLM, from the NetApp Support Site. The RSA enables technical personnel to use the SP or 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 technical support, reducing your involvement in the troubleshooting process. The RSA is not bundled with Data ONTAP. For more information about the RSA, see the *Remote Support Agent Configuration Guide for 7-Mode for Use with Data ONTAP* and the NetApp Remote Support Diagnostics Tool page on the NetApp Support Site.

Related information

NetApp Remote Support Diagnostics Tool page: support.netapp.com/NOW/download/tools/rsa

Managing a system remotely by using the Service Processor

The Service Processor (SP) is a remote management device that enables you to access, monitor, and troubleshoot a system remotely.

The SP provides the following capabilities:

- The SP enables you to access a system remotely to diagnose, shut down, power-cycle, or reboot the system, regardless of the state of the system 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 its 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 serial console and run Data ONTAP commands remotely.

You can access the SP from the serial console or access the serial console from the SP. The SP allows you to open both an SP CLI session and a separate console session simultaneously.

For instance, when a temperature sensor becomes critically high or low, Data ONTAP triggers the SP to shut down the motherboard gracefully. The serial console becomes unresponsive, but you

can still press Ctrl-G on the console to access the SP CLI. You can then use the `system power on` or `system power cycle` command from the SP to power on or power-cycle the system.

- The SP monitors environmental sensors and logs events to help you take timely and effective service actions.

The SP monitors the system temperatures, voltages, currents, and fan speeds. When an environmental sensor has reached an abnormal condition, the SP logs the abnormal readings, notifies Data ONTAP of the issue, and sends alerts and “down system” notifications as necessary through an AutoSupport message, regardless of whether the 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 AutoSupport functionality. The AutoSupport configuration settings and message content behavior are inherited from Data ONTAP.

Note: The SP does not rely on the `autosupport.support.transport` option to send notifications. The SP uses the Simple Mail Transport Protocol (SMTP).

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 events such as boot progress, Field Replaceable Unit (FRU) changes, Data ONTAP-generated events, and SP command history.

- The SP has a nonvolatile memory buffer that stores up to 4,000 events in a system event log (SEL) to help you diagnose issues.

The SEL stores each audit log entry as an audit event. It is stored in onboard flash memory on the SP. 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:

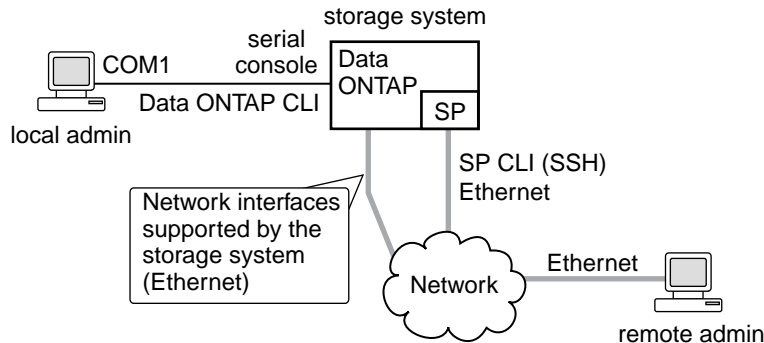
- Hardware events detected by the SP—for example, sensor status about power supplies, voltage, or other components
- Errors detected by the SP—for example, a communication error, a fan failure, or a memory or CPU error
- Critical software events sent to the SP by the system—for example, a panic, a communication failure, a boot failure, or a user-triggered “down system” as a result of issuing the `SP system reset` or `system power cycle` command
- The SP monitors the serial console regardless of whether administrators are logged in or connected to the console.

When messages are sent to the console, the SP stores them in the console log. The console log persists as long as the SP has power from either of the system power supplies. Because the SP operates with standby power, it remains available even when the system is power-cycled or turned off.

- Hardware-assisted takeover is available if the SP is configured.

For more information about hardware-assisted takeover, see the *Data ONTAP High Availability and MetroCluster Configuration Guide for 7-Mode*.

The following diagram illustrates access to the storage system and the SP:



Related concepts

[What the e0M interface is](#) on page 31

Configuring the SP network

Before you can access the SP, the SP network must be configured and enabled. You can configure the SP to use IPv4, IPv6, or both. The SP IPv4 configuration supports static and DHCP addressing, and the SP IPv6 configuration supports static addressing only.

Before you begin

To configure IPv6 connections for the SP, IPv6 must already be configured and enabled for Data ONTAP. The `ip.v6.enable` option controls the IPv6 settings for Data ONTAP. For more information about IPv6 configuration, see the *Data ONTAP Network Management Guide for 7-Mode*.

Steps

1. Configure and enable the SP by using the `system node service-processor network modify` command.
 - The `-address-type` parameter specifies whether the IPv4 or IPv6 configuration of the SP is to be modified.
 - The `-enable` parameter enables the network interface of the specified IP address type.
 - The `-dhcp` parameter specifies whether to use the network configuration from the DHCP server or the network address that you provide.
You can enable DHCP (by setting `-dhcp` to `v4`) only if you are using IPv4. You cannot enable DHCP for IPv6 configurations.
 - The `-ip-address` parameter specifies the public IP address for the SP.
 - The `-netmask` parameter specifies the netmask for the SP (if using IPv4.)
 - The `-prefix-length` parameter specifies the network prefix-length of the subnet mask for the SP (if using IPv6.)
 - The `-gateway` specifies the gateway IP address for the SP.

2. Display the SP network configuration to verify the settings by using the `system node service-processor network show` command.

Example of configuring the SP network

The following example configures the SP to use IPv4, enables the SP, and displays the SP network configuration to verify the settings.

```
system> system node service-processor network modify -node local
-address-type IPv4 -enable true -ip-address 192.168.123.98
-netmask 255.255.255.0 -gateway 192.168.123.1

system> system node service-processor network show -instance -node local

                Node: system
            Address Type: IPv4
        Interface Enabled: true
            Type of Device: SP
                Status: online
            Link Status: up
            DHCP Status: none
                IP Address: 192.168.123.98
            MAC Address: ab:cd:ef:fe:ed:02
                Netmask: 255.255.255.0
Prefix Length of Subnet Mask: -
Router Assigned IP Address: -
    Link Local IP Address: -
        Gateway IP Address: 192.168.123.1

                Node: system
            Address Type: IPv6
        Interface Enabled: false
            Type of Device: SP
                Status: online
            Link Status: disabled
            DHCP Status: none
                IP Address: -
            MAC Address: ab:cd:ef:fe:ed:02
                Netmask: -
Prefix Length of Subnet Mask: -
Router Assigned IP Address: -
    Link Local IP Address: -
        Gateway IP Address: -
2 entries were displayed.

system>
```

Related concepts

Managing AutoSupport on page 157

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 Data ONTAP root password when you use the SP naroot account 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`

Accessing the SP from an administration host

You can log in to the SP from an administration host to perform system management tasks remotely.

Before you begin

The following conditions must be met:

- The administration host you use to access the SP must support SSHv2.
- Your user account must already be set up for accessing the SP.

The accounts that can access the SP include the predefined “naroot” account and Data ONTAP user accounts 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 an IPv4 or IPv6 address, and if five SSH login attempts from a host fail consecutively within 10 minutes, the SP rejects SSH login requests and suspends the communication with the IP address of the host for 15 minutes. The communication resumes after 15 minutes, and you can try to log in to the SP again.

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.

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.

For security reasons, the SP prevents you from logging in with the Data ONTAP “root” account. Instead, it maps the Data ONTAP root account to the SP naroot account. If you use the SP naroot account to access the SP, you also use the Data ONTAP root password. Disabling the Data ONTAP root account also disables the SP naroot account.

The following reserved names cannot be used as account names for accessing the SP—“adm”, “bin”, “cli”, “daemon”, “ftp”, “games”, “halt”, “lp”, “mail”, “man”, “netapp”, “news”, “nobody”, “operator”, “shutdown”, “sshd”, “sync”, “sys”, “uucp”, and “www”.

Steps

1. Enter the following command from the administration host to log in to the SP:

```
ssh username@SP_IP_address
```

2. When 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:
SP>
```

The following example shows how to log in to the SP with a user account, joe, which has been set up to access the SP.

```
[admin_host]$ ssh joe@192.168.123.98
joe@192.168.123.98's password:
SP>
```

The following examples show how to use the IPv6 global address or IPv6 router-advertised address to log in to the SP on a system that has SSH set up for IPv6 and the SP configured for IPv6.

```
[admin_host]$ ssh naroot@fd22:8b1e:b255:202::1234
naroot@fd22:8b1e:b255:202::1234's password:
SP>
```

```
[admin_host]$ ssh naroot@fd22:8b1e:b255:202:2a0:98ff:fe01:7d5b
naroot@fd22:8b1e:b255:202:2a0:98ff:fe01:7d5b's password:
SP>
```

Accessing the SP from the serial console

You can access the SP from the serial console to perform monitoring or troubleshooting tasks.

Steps

1. To access the SP CLI from the serial console, press Ctrl-G at the prompt.
2. Log in to the SP CLI when you are prompted.

The SP prompt appears, indicating that you have access to the SP CLI.

3. To exit the SP CLI and return to the serial console, press Ctrl-D and then press Enter.

Example of accessing the SP CLI from the serial console

The following example shows the result of pressing Ctrl-G from the serial console to access the SP CLI. The `help system power` command is entered at the SP prompt, followed by pressing Ctrl-D and then Enter to return to the serial console.

```
system>
```

(Press Ctrl-G to access the SP CLI.)

```
Switching console to Service Processor
Service Processor Login:
Password:
SP>
SP> 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>
```

(Press Ctrl-D and then Enter to return to the serial console.)

```
system>
```

Accessing the serial console from the SP

The SP's `system console` command enables you to log in to the serial console 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.
3. To exit from the serial console and return to the SP CLI, press Ctrl-D.

Example of accessing the serial 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 console, followed by pressing Ctrl-D, which returns you to the SP prompt.

```
SP> system console
Type Ctrl-D to exit.
```

(Press Enter to see the storage system prompt.)

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

(Command output is displayed.)

(Press Ctrl-D to exit the serial console and return to the SP CLI.)

```
SP>
```

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 press 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.

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 Configuration Guide for 7-Mode for Use with Data ONTAP*.

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
```

Commands for managing a system at the SP admin privilege level

The SP commands at the admin privilege level enable you to display events, logs, and status information for system power, batteries, sensors, field-replaceable units (FRUs), or the SP itself. The commands also enable you to reboot the system or the SP and create a core dump.

The following SP commands are available at the admin privilege level:

Note: Some commands are platform-specific and might not be available on your platform.

If you want to...	Use this command...
Display system date and time	<code>date</code>
Display events that are logged by the SP	<code>events {all info newest <i>number</i> oldest <i>number</i> search <i>keyword</i>}</code>
Exit the SP CLI	<code>exit</code>
Display a list of available commands or subcommands of a specified command	<code>help [<i>command</i>]</code>
Set the privilege level to access the specified mode for the SP CLI	<code>priv set {admin advanced diag}</code> Attention: You should use advanced or diag commands only under the guidance of technical support.

If you want to...	Use this command...
Display the current privilege level for the SP CLI	<code>priv show</code>
Manage the Remote Support Agent (RSA) if it is installed on the system	<code>rsa</code> Note: For information about the RSA, see the <i>Remote Support Agent Configuration Guide for 7-Mode for Use with Data ONTAP</i> .
Display the SP log archives or the files in an archive	<code>sp log history show [-archive {latest all archive-name}] [-dump {all file-name}]</code>
Reboot the SP	<code>sp reboot</code>
Display SP status and network configuration information	<code>sp status [-v -d]</code> Note: The <code>-v</code> option displays SP statistics in verbose form. The <code>-d</code> option adds the SP debug log to the display. The Data ONTAP <code>sysconfig</code> command displays the status for both the storage system and the SP.
Update the SP firmware by using the image at the specified location	<code>sp update image_URL</code> Note: <code>image_URL</code> must not exceed 200 characters.
Display 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>sp uptime</code>
Display ACP information or the status for expander sensors	<code>system acp [show sensors show]</code>
Display battery information	<code>system battery show</code>
Log in to the system console	<code>system console</code> Note: You use Ctrl-D to exit from the system console and return to the SP CLI.

If you want to...	Use this command...
Create a core dump and reset the system	<pre>system core</pre> <p>Note: This command has the same effect as pressing the Non-maskable Interrupt (NMI) button on a system. The SP stays operational as long as the input power to the system is not interrupted.</p>
Display the settings for collecting system forensics on a watchdog reset event, display system forensics information collected during a watchdog reset event, or clear the collected system forensics information.	<pre>system forensics [show log dump log clear]</pre>
List all system FRUs and their IDs	<pre>system fru list</pre>
Display product information for the specified FRU	<pre>system fru show fru_id</pre> <p>Note: You can display FRU IDs by using the <code>system fru list</code> command.</p>
Display console logs	<pre>system log</pre>
Turn the system on or off, or perform a power-cycle (turning the power off and then back on)	<pre>system power {on off cycle}</pre> <p>Note: 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</code> command to turn off or power-cycle the system might cause an improper shutdown of the system (also called a <i>dirty shutdown</i>) and is not a substitute for a graceful shutdown using the Data ONTAP <code>halt</code> command.</p>
Display the status for the power supply	<pre>system power status</pre>
Reset the system by using the specified BIOS firmware image	<pre>system reset {primary backup current}</pre> <p>Note: The SP stays operational as long as the input power to the system is not interrupted.</p>

If you want to...	Use this command...
Display the status for the environmental sensors, including their states and current values	<code>system sensors</code> Note: This command has an equivalent command, <code>system sensors show</code> .
Display the status and details for the specified sensor	<code>system sensors get sensor_name</code> Note: You can obtain <i>sensor_name</i> by using the <code>system sensors</code> or the <code>system sensors show</code> command.
Display the SP hardware and firmware version information	<code>version</code>

Commands for managing a system at the SP advanced privilege level

You can use the SP advanced privilege level to display the SP command history, SP debug file, SP messages file, and data history for field-replaceable units (FRUs). You can also manage the battery firmware and automatic update.

The following SP commands are available only at the advanced privilege level:

Attention: You should use advanced commands only under the guidance of technical support.

If you want to...	Use this command...
Display the SP command history	<code>sp log audit</code>
Display the SP debug information	<code>sp log debug</code>
Display the SP messages file	<code>sp log messages</code>
Display the status of battery firmware automatic update, or enable or disable battery firmware automatic update upon next SP boot	<code>system battery auto_update [status enable disable]</code>
Update the battery firmware from the image at the specified location	<code>system battery flash image_URL</code> Note: You use this command if the automatic battery firmware upgrade process has failed for some reason.
Compare the current battery firmware image against a specified firmware image	<code>system battery verify [image_URL]</code> Note: If <i>image_URL</i> is not specified, the default battery firmware image is used for comparison.
Display the FRU data history log	<code>system fru log show</code>

Related tasks

[Setting the privilege level](#) on page 26

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 in the SP CLI:

```
SP toaster> 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
SBTY_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` in the SP CLI for the threshold-based sensor 5V:

```
SP toaster> 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 in the SL CLI 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 in the SL CLI.

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, 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 toaster> 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 toaster> 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]
```

Troubleshooting a system by using the SP

When you encounter a problem with the storage system, you can use the SP to display information about the problem, create a core dump, and reboot the system, even if the system's firmware is corrupted.

The following table describes the common SP commands that you can use at the SP prompt to troubleshoot a system:

If this condition occurs...	And you want to...	Enter this command at the SP CLI prompt...
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>

If this condition occurs...	And you want to...	Enter this command at the SP CLI prompt...
The system is not responding properly.	Access the system console from the SP.	<code>system console</code>
	Create a core dump and reboot the system.	<code>system core</code>
	Power-cycle the system.	<code>system power cycle</code>
You receive an AutoSupport message indicating an event such as a panic or hardware component failure.	Display what has occurred at the 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 system firmware is corrupted.	Boot the system by using the backup image of the 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

[Commands for managing a system at the SP admin privilege level](#) on page 198

[Commands for managing a system at the SP advanced privilege level](#) on page 201

Managing the SP with Data ONTAP

You can use Data ONTAP to set up and display the SP configuration, display the SP status, reboot the SP, manage the SP firmware image, and manage access to the SP.

Methods of managing SP firmware updates

Starting with Data ONTAP 8.2, a baseline SP firmware image is packaged with the Data ONTAP image. By default, the SP automatic update functionality is enabled. You have the option to manually trigger an SP update.

Data ONTAP 8.2 and later releases include an SP firmware image that is called the *baseline image*. You do not need to download the baseline SP firmware image separately. If a new version of the SP firmware becomes subsequently available, you have the option to download it from the [System](#)

[Firmware and Diagnostics Download](#) page on the NetApp Support Site and update the SP firmware to the downloaded version without upgrading the Data ONTAP version. For information about manually downloading and updating the SP firmware, see the SP Firmware Download and Installation Instructions on the NetApp Support Site.

Data ONTAP offers the following methods for managing SP firmware updates:

- The SP automatic update functionality is enabled by default, allowing the SP firmware to be automatically updated in the following scenarios:
 - When you upgrade to a new version of Data ONTAP
The Data ONTAP upgrade process automatically includes the SP firmware update, provided that the SP firmware version bundled with Data ONTAP is newer than the SP version running on the system.

Note: Data ONTAP detects a failed SP automatic update and triggers a corrective action to retry the SP automatic update up to three times. If all three retries have failed, you should contact technical support.

- When you download a version of the SP firmware from the NetApp Support Site and the downloaded version is newer than the one that the SP is currently running

You have the option to disable the SP automatic update functionality by using the `system node service-processor image modify` command. However, it is best to leave the functionality enabled. Disabling the functionality can result in suboptimal or nonqualified combinations between the Data ONTAP image and the SP firmware image.

- Data ONTAP enables you to trigger an SP update manually and specify how the update should take place by using the `system node service-processor image update` command.

You can specify the following options:

- The SP firmware package to use (`-package`)
You can update the SP firmware to a downloaded package by specifying the package file name. The `system node image package show` command displays all package files (including the files for the SP firmware package) that are available on a node.
- Whether to use the baseline SP firmware package for the SP update (`-baseline`)
You can update the SP firmware to the baseline version that is bundled with the currently running version of Data ONTAP.
- Whether to update the entire firmware image or only the changed portions (`-update-type`)
- If updating the entire firmware image, whether to also reset log settings to the factory default and clear contents of all logs maintained by the SP, including the event logs, IPMI logs, and forensics logs (`-clear-logs`)

For information about the `system node service-processor image update` command, see the man page.

- Data ONTAP enables you to display the status for the latest SP firmware update by using the `system node service-processor image update-progress show` command.

Related information

NetApp Support Site: support.netapp.com

Restricting SP access to only the specified administration hosts

You can configure the SP to accept SSH requests from only the administration hosts that you specify.

Step

1. Enter the following command to specify the administration host or hosts that you want to grant SP access:

```
options sp.ssh.access host_spec
```

You can specify *host_spec* in the following forms:

- `host[=|!=]host_list`
host_list is a comma-separated list that includes host names, IP addresses, or IP addresses with a netmask.
- `all` or `*`
Allows all hosts to access the SP.
- `none` or `-`
Allows no hosts to access the SP.

The default for *host_spec* is `*`.

For more information and examples about using this option, see the `na_spaccess(8)` man page.

Examples of restricting SP access to only the specified hosts

The following example grants SP SSH access to the administration host with the specified IP address:

```
system> options sp.ssh.access host=192.168.123.98
```

The following example grants SP SSH access to two administration hosts, identified by their host names:

```
system> options sp.ssh.access host=myhost1,myhost2
```

The following example grants SP SSH access to all hosts with their IP address prefix matching 3FFE:81D0:107:2082:

```
system> options sp.ssh.access host=3FFE:81D0:107:2082::1/64
```

Configuring automatic logout of idle SSH connections to the SP

You can configure the automatic logout settings so that an SSH connection to the SP is automatically terminated after the connection has been idle for the number of minutes you specify.

About this task

Setting changes for automatic logout of idle SP SSH connections take effect only on SSH sessions that start after the changes.

Automatic logout does not take effect if you access the SP through the serial console.

Steps

1. Enter the following command to enable SSH automatic logout for the SP:

```
options sp.autologout.enable on
```

Note: The default is `on`. Setting the option to `off` disables SSH automatic logout for the SP, causing the `sp.autologout.timeout` option to have no effect.

2. Enter the following command to specify the number of minutes after which an idle SSH connection to the SP is automatically disconnected:

```
options sp.autologout.timeout minutes
```

The default is 60 minutes.

Example of configuring automatic logout of idle SSH connections to the SP

The following example configures the SP to automatically disconnect SSH sessions that are idle for 30 minutes or more:

```
system> options sp.autologout.enable on
system> options sp.autologout.timeout 30
```

Data ONTAP commands for managing the SP

Data ONTAP provides commands for managing the SP, including setting up and displaying the SP network configuration, displaying the current SP status, rebooting the SP, managing the SP firmware image, and managing SSH access to the SP.

You can use the following Data ONTAP commands and options to manage the SP:

If you want to...	Use this Data ONTAP command...
<p>Set up or modify the SP network configuration, including the following:</p> <ul style="list-style-type: none"> • The IP address type (IPv4 or IPv6) • Whether the network interface of the specified IP address type should be enabled • If you are using IPv4, whether to use the network configuration from the DHCP server or the network address that you specify • The public IP address for the SP • The netmask for the SP (if using IPv4) • The network prefix-length of the subnet mask for the SP (if using IPv6) • The gateway IP address for the SP 	<pre>system node service-processor network modify</pre>
<p>Display the SP network configuration, including the following:</p> <ul style="list-style-type: none"> • The configured address type (IPv4 or IPv6) and whether it is enabled • The remote management device type • The current SP status and link status • Network configuration, such as IP address, MAC address, netmask, prefix-length of subnet mask, router-assigned IP address, link local IP address, and gateway IP address 	<pre>system node service-processor network show</pre> <p>Note: Displaying complete SP network details requires the <code>-instance</code> parameter.</p>
<p>Display general SP information, including the following:</p> <ul style="list-style-type: none"> • The remote management device type • The current SP status • Whether the SP network is configured • Network information, such as the public IP address and the MAC address • The SP firmware version and Intelligent Platform Management Interface (IPMI) version • Whether the SP firmware automatic update is enabled 	<pre>system node service-processor show</pre>

If you want to...	Use this Data ONTAP command...
<p>Reboot the SP and optionally specify the SP firmware image (primary or backup) to use</p>	<pre>system node service-processor reboot-sp</pre> <p>Attention: You should avoid booting the SP from the backup image. Booting from the backup image is reserved for troubleshooting and recovery purposes only. It might require that the SP automatic firmware update be disabled, which is not a recommended setting. You should contact Technical Support before attempting to boot the SP from the backup image.</p>
<p>Display the details of the currently installed SP firmware image, including the following:</p> <ul style="list-style-type: none"> • The remote management device type • The partition (primary or backup) that the SP is booted from, its status, and firmware version • Whether the firmware automatic update is enabled and the last update status 	<pre>system node service-processor image show</pre> <p>Note: The <code>-is-current</code> parameter indicates the partition (primary or backup) that the SP is currently booted from, not whether the installed firmware version is most current.</p>
<p>Enable or disable the SP automatic firmware update</p>	<pre>system node service-processor image modify</pre> <p>Note: By default, the SP firmware is automatically updated with the update of Data ONTAP or when a new version of the SP firmware is manually downloaded. Disabling the automatic update is not recommended because doing so can result in suboptimal or nonqualified combinations between the Data ONTAP image and the SP firmware image.</p>
<p>Manually download an SP firmware image</p>	<pre>software get</pre> <p>Note: The SP firmware image is packaged with Data ONTAP. You do not need to download the SP firmware manually, unless you want to use an SP firmware version that is different from the one packaged with Data ONTAP.</p>

If you want to...	Use this Data ONTAP command...
<p>Manually update the SP firmware, by specifying the following:</p> <ul style="list-style-type: none"> • The SP firmware package to use You can have the SP use a specific SP firmware package by specifying the package file name. The <code>software list</code> command displays all package files (including the files for the SP firmware package) that are available on a node. • The installation baseline You can update the SP firmware to the baseline version that is bundled with the currently running version of Data ONTAP. • Whether to update the entire firmware image or only the changed portions • If updating the entire firmware image, whether to also reset log settings to the factory default and clear contents of all logs maintained by the SP, including the event logs, IPMI logs, and forensics logs 	<pre>system node service-processor image update</pre>
<p>Display the status for the latest SP firmware update, including the following information:</p> <ul style="list-style-type: none"> • The start and end time for the latest SP firmware update • Whether an update is in progress and the percentage that is complete 	<pre>system node service-processor image update-progress show</pre>
<p>Enable or disable automatic logout of idle SSH connections to the SP</p>	<pre>options sp.autologout.enable</pre>
<p>Specify the number of minutes after which an idle SSH connection to the SP is automatically disconnected</p>	<pre>options sp.autologout.timeout</pre> <p>Note: For this option to take effect, the <code>sp.autologout.enable</code> option must be set to on.</p>
<p>Restrict SP access to only the specified administration hosts</p>	<pre>options sp.ssh.access</pre>

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.

About this task

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 214

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 214

Managing a system remotely by using the Remote LAN Module

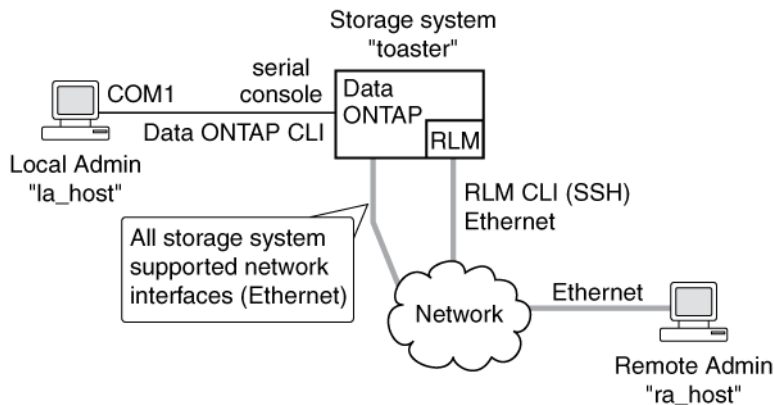
The Remote LAN Module (RLM) is a remote management device that is provided on the 31xx, 6040, and 6080 platforms. The RLM provides remote system management capabilities, including remote access, monitoring, troubleshooting, logging, and alerting features.

The RLM stays operational regardless of the operating state of the system. It is powered by a standby voltage, which is available as long as the system has input power to at least one of its 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 system issues. For example, it might help a remote administrator determine if the system was shut down due to an extreme temperature change.

For instructions about how to cable a 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 access the system through the serial console or from an Ethernet connection using any supported network interface.
You use the Data ONTAP CLI to administer the system.
- With the RLM, you can *remotely* access the system through the serial console.
The RLM is directly connected to the system through the serial console. You use the Data ONTAP CLI to administer the system and the RLM.
- With the RLM, you can also access the system through an Ethernet connection using a secure shell client application.
You use the RLM CLI to monitor and troubleshoot the 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.

Related concepts

What the e0M interface is on page 31

What the RLM does

The commands in the RLM CLI 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.

Note: The RLM does not rely on the `autosupport.support.transport` option to send notifications. The RLM uses the Simple Mail Transport Protocol (SMTP).

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.

The RLM has a nonvolatile memory buffer that stores up to 4,000 system events in a system event log (SEL) 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.

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

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 High Availability and MetroCluster Configuration Guide for 7-Mode*.

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.

Related concepts

[Troubleshooting the storage system by using the RLM](#) on page 231

[Managing AutoSupport](#) on page 157

Ways to configure the RLM

Before using the RLM, you must configure it for your system and network. You can configure the RLM when setting up a new system with the RLM already installed, after setting up a new system with the RLM already installed, or when adding an RLM to an existing system.

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

- Initializing a storage system that has the RLM preinstalled
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 Software Setup Guide for 7-Mode*.
- 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.

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 an IPv4 address for the RLM, you need the following information:

- An available static IP address
- The netmask of your network
- The gateway of your network

If you are using IPv6 for RLM static addressing, you need the following information:

- The IPv6 global address
- The subnet prefix for the RLM
- The IPv6 gateway for the RLM
- AutoSupport information

The RLM sends event notifications based on the following AutoSupport settings:

- `autosupport.to`
- `autosupport.mailhost`

It is best that you configure at least the `autosupport.to` option 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 to the system administrative recipients specified in the `autosupport.to` option. You are prompted to enter the name or the IP address of the AutoSupport mail host when you configure the RLM.

Related concepts

[*Managing AutoSupport*](#) on page 157

Configuring the RLM

You can use the `setup` command or the `rlm setup` command to configure the RLM. You can configure the RLM to use either a static or a DHCP address.

Before you begin

AutoSupport should be configured before you configure the RLM. Data ONTAP automatically sends the AutoSupport configuration to the RLM, allowing the RLM to send alerts and notifications through AutoSupport messages.

About this task

If you are running RLM firmware version 4.0 or later, and you have enabled IPv6 for Data ONTAP, you have the option to configure the RLM for only IPv4, for only IPv6, or for both IPv4 and IPv6. Disabling IPv6 on Data ONTAP also disables IPv6 on the RLM.

Attention: If you disable both IPv4 and IPv6, and if DHCP is also not 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`.

Note: DHCPv6 servers are not currently supported.

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

Note: Entering `0.0.0.0` for the static IP address disables IPv4 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
5. If you enabled IPv6 for Data ONTAP and your RLM firmware version is 4.0 or later, the RLM supports IPv6, and the RLM setup asks you whether to configure IPv6 connections for the RLM:
 - To configure IPv6 connections for the RLM, enter `y`.
 - To disable IPv6 connections for the RLM, enter `n`.

Note: You can use the `rlm status` command to find the RLM version information.

6. If you choose to configure IPv6 for the RLM, provide the following IPv6 information when prompted by the RLM setup:
 - The IPv6 global address

Even if no IPv6 global address is assigned for the RLM, the link-local address is present on the RLM. The IPv6 router-advertised address is also present if the `ip.v6.ra_enable` option is set to `on`.

- The subnet prefix for the RLM
- The IPv6 gateway for the RLM

Note: You cannot use the RLM setup to enable or disable the IPv6 router-advertised address for the RLM. However, when you use the `ip.v6.ra_enable` option to enable or disable the IPv6 router-advertised address for Data ONTAP, the same configuration applies to the RLM.

For information about enabling IPv6 for Data ONTAP or information about global, link-local, and router-advertised addresses, see the *Data ONTAP Network Management Guide for 7-Mode*.

7. At the storage system prompt, enter the following command to verify that the RLM network configuration is correct:

```
rlm status
```

8. 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 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 for both IPv4 and IPv6 connections:

```
toaster> 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 event 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
Do you want to enable IPv6 on the RLM ? y
Do you want to assign IPv6 global address? y
Please enter the IPv6 address for the RLM []:fd22:8b1e:b255:204::1234
Please enter the subnet prefix for the RLM []: 64
Please enter the IPv6 Gateway for the RLM []:fd22:81be:b255:204::1
Verifying mailhost settings for RLM use...
```

The following example shows that the RLM is configured to use DHCP and IPv6:

```
toaster> 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
Do you want to enable IPv6 on the RLM ? y
Do you want to assign IPv6 global address? y
Please enter the IPv6 address for the RLM [fd22:8b1e:b255:204::1234]:
Please enter the subnet prefix for the RLM [64]:
Please enter the IPv6 Gateway for the RLM [fd22:81be:b255:204::1]:
Verifying mailhost settings for RLM use...

```

The following example displays the RLM status and configuration information:

```

toaster> rlm status
  Remote LAN Module      Status: Online
    Part Number:         110-00030
    Revision:            A0
    Serial Number:       123456
    Firmware Version:    4.2
    Mgmt MAC Address:    00:A0:98:01:7D:5B
    Ethernet Link:       up, 100Mb, full duplex, auto-neg complete
    Using DHCP:          no
IPv4 configuration:
  IP Address:            192.168.123.98
  Netmask:               255.255.255.0
  Gateway:               192.168.123.1
IPv6 configuration:
  Global IP:             fd22:8b1e:b255:204::1234
  Prefix Length:         64
  Gateway:               fd22:81be:b255:204::1
  Router Assigned IP:    fd22:8b1e:b255:204:2a0:98ff:fe01:7d5b
  Prefix Length:         64
  Link Local IP:         fe80::2a0:98ff:fe00:7d1b
  Prefix Length:         64

```

Related concepts

[Prerequisites for configuring the RLM](#) on page 218

[Managing AutoSupport](#) on page 157

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 Data ONTAP root password when you use the RLM's naroot account 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 95

[*Predefined roles*](#) on page 106

[*Supported capability types*](#) on page 108

Related tasks

[*Creating a new role and assigning capabilities to roles*](#) on page 110

[*Modifying an existing role or its capabilities*](#) on page 111

[*Disabling root account access to the storage system*](#) on page 98

Restricting RLM access to only the specified administration hosts

You can configure the RLM to accept SSH requests from only the administration hosts that you specify.

Before you begin

Your system must be running RLM firmware 4.1 or later for the RLM access control to be supported. For information about downloading and updating the RLM firmware, see the *Data ONTAP Upgrade and Revert/Downgrade Guide for 7-Mode*.

Step

1. Enter the following command to specify the administration host or hosts that you want to grant RLM access:

```
options rlm.ssh.access host_spec
```

You can specify *host_spec* in the following forms:

- `host[=|!=]host_list`
host_list is a comma-separated list that includes host names, IP addresses, or IP addresses with a netmask.
- `all` or `*`
Allows all hosts to access the RLM.
- `none` or `-`
Allows no hosts to access the RLM.

The default for *host_spec* is `*`.

For more information and examples about using this option, see the `na_rlmaccess(8)` man page.

Examples of restricting RLM access to only the specified hosts

The following example grants RLM SSH access to the administration host with the specified IP address:

```
system> options rlm.ssh.access host=192.168.123.98
```

The following example grants RLM SSH access to two administration hosts, identified by their host names:

```
system> options rlm.ssh.access host=myhost1,myhost2
```

The following example grants RLM SSH access to all hosts with their IP address prefix matching 3FFE:81D0:107:2082:

```
system> options rlm.ssh.access host=3FFE:81D0:107:2082::1/64
```

Configuring automatic logout of idle SSH connections to the RLM

You can configure the automatic logout settings so that an SSH connection to the RLM is automatically terminated after the connection has been idle for the number of minutes you specify.

Before you begin

Your system must be running RLM firmware version 4.1 or later for the automatic logout configuration to be supported. For information about downloading and updating the RLM firmware, see the *Data ONTAP Upgrade and Revert/Downgrade Guide for 7-Mode*.

About this task

Setting changes for automatic logout of idle RLM SSH connections take effect only on SSH sessions that start after the changes.

Automatic logout does not take effect if you access the RLM through the serial console.

Steps

1. Enter the following command to enable SSH automatic logout for the RLM:

```
options rlm.autologout.enable on
```

Note: The default is `on`. Setting the option to `off` disables SSH automatic logout for the RLM, causing the `rlm.autologout.timeout` option to have no effect.

2. Enter the following command to specify the number of minutes after which an idle SSH connection to the RLM is automatically disconnected:

```
options rlm.autologout.timeout minutes
```

The default is 60 minutes.

Example of configuring automatic logout of idle SSH connections to the RLM

The following example configures the RLM to automatically disconnect SSH sessions that are idle for 30 minutes or more:


```
system> options rlm.autologout.enable on
system> options rlm.autologout.timeout 30
```

Logging in to the RLM from an administration host

You can log in to the RLM from an administration host to perform administrative tasks remotely, if the host has a Secure Shell client application that supports SSHv2 and you have proper administrative privileges on the storage system.

Before you begin

The following are the prerequisites for logging in to the RLM:

- A secure shell application must be installed on the host.
The RLM accepts only SSH connections. It does not respond to other protocols.
RLM firmware version 4.0 or later accepts only SSHv2 access to the RLM.
- You must have access to 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

About this task

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.

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.

Steps

1. Enter the following command from the UNIX host:

```
ssh username@RLM_IP_address
```

2. When you are prompted, enter the password for *username*.

The RLM prompt appears, indicating that you have access to the RLM CLI.

Examples of RLM access from an administration 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 to access the RLM.

```
ssh joe@192.168.123.98
```

The following examples show how to use the IPv6 global address or IPv6 router-advertised address to log in to the RLM on a storage system that has SSH set up for IPv6 and the RLM configured for IPv6.

```
ssh joe@fd22:8b1e:b255:202::1234
```

```
ssh joe@fd22:8b1e:b255:202:2a0:98ff:fe01:7d5b
```

Accessing the serial console from the RLM

The RLM's `system console` command enables you to log in to the serial console 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 system prompt.
3. To exit the serial console and return to the RLM CLI, press Ctrl-D.

Example of accessing the serial 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 console, followed by Ctrl-D, which returns you to the RLM prompt.

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

(Press Enter to see the storage system prompt.)

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

(Command output is displayed.)

(Press Ctrl-D to exit the storage serial console and return to the RLM CLI.)

```
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 press 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.

If you use the RLM to 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, you can keep the RLM CLI session and the system console session active simultaneously. The system console session provides real-time output from the system, including the progress of the system boot.

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 Configuration Guide for 7-Mode for Use with Data ONTAP*.

- 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
```

Commands for managing the RLM at the admin privilege level

You can perform most RLM tasks at the admin privilege level. For example, you can display system events and status information for environmental sensors, reboot the storage system or the RLM, and create a system core dump.

The following RLM commands are available at the admin privilege level:

If you want to...	Use this command...
Display system date and time	<code>date</code>
Display storage system events logged by the RLM	<code>events {all info newest oldest search string }</code>
Exit the RLM CLI	<code>exit</code>
Display a list of available commands or subcommands of a specified command	<code>help [command]</code>

If you want to...	Use this command...
Set the privilege level to access the specified mode	<code>priv set {admin advanced diag}</code>
Display the current privilege level	<code>priv show</code>
Reboot the RLM	<code>rlm reboot</code>
Display the RLM environmental sensor status	<code>rlm sensors [-c]</code> Note: The <code>-c</code> option, which takes a few seconds to display, shows current values rather than cached values.
Display RLM status	<code>rlm status [-v -d]</code> Note: The <code>-v</code> option displays verbose statistics. The <code>-d</code> option displays RLM debug information. The Data ONTAP <code>sysconfig</code> command displays both the storage system and RLM status.
Update the RLM firmware	<code>rlm update http://path [-f]</code> Note: The <code>-f</code> option issues a full image update.
Manage the RSA if it is installed on your storage system	<code>rsa</code> Note: For information about the RSA, see the <i>Remote Support Agent Configuration Guide for 7-Mode for Use with Data ONTAP</i> .
Log in to the Data ONTAP CLI	<code>system console</code> Note: Pressing Ctrl-d returns you to the RLM CLI.
Dump the system core and reset the system	<code>system core</code> Note: This command has the same effect as pressing the Non-maskable Interrupt (NMI) button on a storage system. The RLM stays operational as long as input power to the storage system is not interrupted.

If you want to...	Use this command...
Turn on or turn off the storage system, or perform a power-cycle (which turns off system power and then turns it back on)	<pre>system power {on off cycle}</pre> <p>Note: Standby power stays on, even when the storage system is off. During power-cycling, a brief pause occurs before power is turned back on.</p> <p>Attention: Using the <code>system power</code> command to turn off or power-cycle the storage system might cause an improper shutdown of the system (also called a <i>dirty shutdown</i>) and is not a substitute for a graceful shutdown using the Data ONTAP <code>halt</code> command.</p>
Display status for each power supply, such as presence, input power, and output power	<pre>system power status</pre>
Reset the storage system using the specified BIOS firmware image	<pre>system reset {primary backup current}</pre> <p>Note: The RLM stays operational as long as input power to the storage system is not interrupted.</p>
Display the RLM version information, including hardware and firmware information	<pre>version</pre>

Commands for managing the RLM at the advanced privilege level

In addition to using the RLM admin commands, you can use the RLM advanced privilege level to display RLM command history, RLM debug and message files, status of environmental sensors, and RLM statistics.

The following RLM commands are available only at the advanced privilege level:

If you want to display...	Use this command...
RLM command history or search for audit logs from the system event log (SEL)	<pre>rlm log audit</pre>
RLM debug file	<pre>rlm log debug</pre>
RLM message file	<pre>rlm log messages</pre>
List of environmental sensors, their states, and their current values	<pre>system sensors</pre>

If you want to display...	Use this command...
RLM statistics	<code>rlm status -v</code>

Related tasks

[Setting the privilege level](#) on page 26

Troubleshooting the storage system by using the RLM

When you encounter a problem with the storage system, you can use the RLM to display information about the problem, create a core dump, and reboot the system, even if the system's firmware is corrupted.

The following table describes the RLM commands that you can use to troubleshoot a system:

If this condition occurs...	And you want to...	Enter this command at the RLM CLI prompt...
You receive an AutoSupport message indicating an event such as a panic or hardware component failure.	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 system is not responding properly.	Access the system console from the RLM.	<code>system console</code>
	Create a core dump and reboot the system.	<code>system core</code>
	Power-cycle the system.	<code>system power cycle</code>
The system firmware is corrupted.	Boot the system by using a backup copy of the system firmware.	<code>system reset backup</code>

Managing the RLM with Data ONTAP

You can manage the RLM from Data ONTAP by using the `rlm` commands.

Data ONTAP commands for managing the RLM

Data ONTAP provides the `rlm` commands for managing 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 and options for managing the RLM. These commands are also described in the `na_rlm(1)` man page.

If you want to...	Use this Data ONTAP command...
Initiate the interactive RLM setup script	<code>rlm setup</code>
Display whether the RLM has been configured	<code>options rlm.setup</code>
Display the list of available <code>rlm</code> commands	<code>rlm help</code>
Display the current status of the RLM, including the following: <ul style="list-style-type: none"> Whether the RLM is online The version that the RLM is running Network and configuration information 	<code>rlm status</code>
Reboot the RLM and trigger the RLM to perform a self-test	<code>rlm reboot</code> Note: Any console connection through the RLM is lost during the reboot.
Send a test email to all recipients specified with the <code>autosupport.to</code> option	<code>rlm test autosupport</code> Note: For this command to work, the <code>autosupport.enable</code> and <code>autosupport.mailhost</code> options must be configured properly.
Perform an 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	<code>rlm test snmp</code> Note: For information about SNMP traps, see the <i>Data ONTAP Network Management Guide for 7-Mode</i> .

If you want to...	Use this Data ONTAP command...
Update the RLM firmware	<pre>rlm update</pre> <p>Note: For information about downloading and updating the RLM firmware, see the <i>Data ONTAP Upgrade and Revert/Downgrade Guide for 7-Mode</i>.</p>
Display the RLM update status, including the following: <ul style="list-style-type: none"> • Whether an RLM update is currently in progress • Completion percentage • The start and end time for the update 	<pre>rlm update-status</pre>
Enable or disable automatic logout of idle SSH connections to the RLM	<pre>options rlm.autologout.enable</pre>
Specify the number of minutes after which an idle SSH connection to the RLM is automatically disconnected	<pre>options rlm.autologout.timeout</pre> <p>Note: For this option to take effect, the <code>rlm.autologout.enable</code> option must be set to on.</p>
Restrict RLM access to only the specified administration hosts	<pre>options rlm.ssh.access</pre>

Related concepts

[Ways to configure the RLM](#) on page 217

RLM and SNMP traps

If SNMP is enabled for the RLM, the RLM generates SNMP traps to configured trap hosts for all "down system" 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.

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

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.

About this task

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 about how to download and update the RLM firmware, see the *Data ONTAP Upgrade and Revert/Downgrade Guide for 7-Mode*.

Related tasks

[Disabling SNMP traps for only the RLM](#) on page 234

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 at the storage system prompt:

```
options rlm.snmp.traps off
```

The default 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 234

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.
2. 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
```

3. From the administration host, test the network connection for the RLM by entering the following command:

```
ping rlm_IP_address
```

4. 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 by entering the following command at the storage system prompt:

```
rlm setup
```

- Verify that the administration host has a route to the RLM.
5. Reboot the RLM by entering the following command at the storage system prompt:

```
rlm reboot
```

Note: It takes approximately one minute for the RLM to reboot.

6. If the RLM does not reboot, repeat Steps 2 through 5. If the RLM still does not reboot, contact technical support for assistance.

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.

Displaying 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.
<code>sysconfig -a</code>	Displays the same information as the <code>-v</code> option, but the information is more detailed.
<code>sysconfig -A</code>	<div>Displays storage system information gathered by the following commands, one after the other:</div> <div><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></div> <div>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.</div>

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 -P	Enumerates the PCI hierarchy of the system, enabling the support and debug of errors from the PCI subsystem. The <code>-P</code> option is not supported on all platforms.
sysconfig -p	Displays information about the physical host machine and its mapping to the virtual machine.
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 Network Management Guide for 7-Mode</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>

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 242

Displaying 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:
 - Aggregate name
 - Whether it is an aggregate (32-bit or a 64-bit) or traditional volume
 - Whether it is online, offline, or restricted
 - 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 Storage Management Guide for 7-Mode*. 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 242

Displaying 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:
 - 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 Storage Management Guide for 7-Mode*. 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 242

Commands for displaying environmental status

The `environment` commands enable you to display all environment information, shelf environment status, and chassis environment status.

Data ONTAP runs the `environment` commands 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 can run the `environment` commands manually to monitor the storage system subsystems, especially when you suspect a problem and when reporting abnormal conditions to technical support.

If you want to...	Use this command...
Display all storage system environment information	<code>environment status</code> Note: For systems that contain internal drives, the <code>environment status</code> command displays information for both the internal and the external storage environment.
Display the shelf environmental status	<code>environment status shelf [adapter]</code>
Display the environmental status of chassis components	<code>environment chassis</code>
Display detailed information from all chassis sensors	<code>environment chassis list-sensors</code>

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

Getting Fibre Channel information

You can display Fibre Channel (FC) information such as the link statistics for all disks on a loop, internal FC driver statistics, and the relative physical positions of drives on a loop.

Step

1. To display FC 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 FC driver. The FC 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 FC information, either interactively or with a script, by using the `fcsp` 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 242

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.

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 245

Collecting system information with 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 stats 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 stats 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 stats show command	Add <code>-O name=value</code> to the stats show command. <i>name</i> is the name of the option you want to filter and <i>value</i> is on or off.

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.

You can use the `cifs.perfmon.allowed_users` option to grant `perfmon` access to specified users or groups. By default, no argument is set for the option, and only members of the Administrators group have access to `perfmon`. For more information about the `cifs.perfmon.allowed_users` option, see the `options(1)` man page.

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 tool that reports performance information for both the host and the storage system. You can run it on 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 NetApp Support Site and navigate to **Software Downloads > ToolChest**.

Related information

NetApp Support Site: support.netapp.com

Managing system performance

You can use several features to improve system performance.

Managing storage system resources by using FlexShare

FlexShare enables you to use priorities and hints to increase your control over how your storage system resources are used. This control enables you to consolidate workloads without negatively impacting critical applications.

FlexShare uses the following methods to help you manage your storage system resources:

- 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.

What FlexShare is

FlexShare provides workload prioritization for a storage system. It prioritizes processing resources for key services when the system is under heavy load. FlexShare does not provide guarantees on the availability of resources or how long particular operations will take to complete. FlexShare provides a priority mechanism to give preferential treatment to higher priority tasks.

FlexShare provides the following key features:

- Relative priority of different volumes
- Per-volume user versus system priority
- Per-volume cache policies

FlexShare provides the following key benefits:

- Simplification of storage management
 - Reduces the number of storage systems that need to be managed by enabling consolidation
 - Provides a simple mechanism for managing performance of consolidated environments
- Cost reduction

- Allows increased capacity and processing utilization per storage system without impact to critical applications
- No special hardware or software required
- No additional license required
- Flexibility
 - Can be customized easily to meet performance requirements of different environment workloads

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.

FlexShare is designed to change performance characteristics when the storage system is under load. If the storage system is not under load, it is expected that the FlexShare impact will be minimal and can even be unnoticeable. 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.

Use FlexShare with storage systems that have hard disk drives (HDDs) only. FlexShare is not designed for use with storage systems that have solid-state drives (SSDs). Enabling FlexShare on a storage system that has SSDs can result in decreased throughput to SSD-based volumes.

Related concepts

[FlexShare and the buffer cache policy values](#) on page 255

Related tasks

[Assigning FlexShare priority to a volume relative to other volumes](#) on page 256

[Assigning FlexShare priority to system operations relative to user operations](#) on page 257

FlexShare and priority levels

FlexShare 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 that helps it 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

Considerations for 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. The takeover node creates a new priority policy by merging the policies configured on each individual node. For this reason, you should ensure 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 FlexShare queue works

The default FlexShare queue consists of volumes that have no priority assigned to them. Understanding how the default priority is used helps you create the optimal FlexShare 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.

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. Examples of such volumes include 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.

Using FlexShare

You use FlexShare to assign priorities to volume data access, set the volume buffer cache policy, and modify the default priority.

Assigning FlexShare 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.

About this task

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
```

2. 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.

3. 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 FlexShare queue works](#) on page 255

Assigning FlexShare priority to system operations relative to user operations

If system operations (for example, SnapMirror transfers or `ndmcopy` 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.

About this task

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
```

Using FlexShare to set 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`. This 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 255

Removing FlexShare priority from a volume

You can temporarily disable the FlexShare priority for a specific 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 return to 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 return to 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 FlexShare default priority

If you have not assigned a FlexShare 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.

About this task

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
```

Increasing WAFL cache memory

You can increase Write Anywhere File Layout (WAFL) cache memory in a system that has a caching module installed (Performance Acceleration Module (PAM), Flash Cache module, or Flash Cache 2 module). To increase the WAFL cache memory, you use the WAFL external cache, a software component of Data ONTAP.

WAFL external cache provides extra WAFL cache memory to improve the performance of the storage system by reducing the number of disk reads. You can control how user data blocks are cached by changing the mode of operation for a caching module. You can keep the default mode (normal user data blocks) or you can choose metadata mode or low-priority blocks mode.

You should verify that the WAFL external cache functionality is enabled after you install a caching module.

Note: WAFL external cache does not require a separate license if your system is running Data ONTAP 8.1 or later.

Note: Not all systems have a caching module installed. Therefore, not all systems can utilize the WAFL external cache functionality.

WAFL external cache does not cache data that is stored in a RAID group composed of SSDs.

If you use WAFL external cache on storage systems with a high-availability configuration, you must ensure that the WAFL external cache options are the same on both nodes. Otherwise, a takeover can result in lower performance due to the lack of WAFL external cache on the remaining node.

Besides the Data ONTAP options that you can use to manage WAFL external cache, a diagnostic command is available for sanitizing a caching module. For more information, see the *Diagnostics Guide*.

How Flash Pools and Flash Cache compare

Both the Flash Pool technology and the family of Flash Cache modules (Flash Cache and Flash Cache 2) provide a high-performance cache to increase storage performance. However, there are differences between the two technologies that you should understand before choosing between them.

You can employ both technologies on the same system. However, data stored in volumes associated with a Flash Pool (or an SSD aggregate) is not cached by Flash Cache.

Criteria	Flash Pool	Flash Cache
Scope	A specific aggregate	All aggregates assigned to a controller
Caching types supported	Read and write	Read
Cached data availability during and after takeover events	Cached data is available and unaffected by either planned or unplanned takeover events.	Cached data is not available during takeover events. After giveback for a planned takeover, previously cached data that is still valid is re-cached automatically.
PCIe slot on storage controller required?	No	Yes
Supported with array LUNs?	No	Yes
Supported with Storage Encryption?	No	Yes. Data in the cache is not encrypted.
Supported with SnapLock?	No	Yes

For more information about Flash Pools, see the *Data ONTAP Storage Management Guide for 7-Mode*.

Enabling and disabling WAFL external cache

You can enable or disable the WAFL external cache functionality for a storage system that has a caching module installed (Performance Acceleration Module, Flash Cache module, or Flash Cache 2 module). You should verify that the WAFL external cache functionality is enabled after you install a caching module.

About this task

The `flexscale.enable` option enables or disables the WAFL external cache functionality. If your storage system does not have a caching module installed, the `flexscale.enable` option enables or disables the Predictive Cache Statistics (PCS). PCS is supported on platforms that support caching modules.

WAFL external cache does not require a separate license if your system is running Data ONTAP 8.1 or later. PCS does not require a license.

Steps

1. To verify whether the WAFL external cache is enabled or disabled, enter the following command:

```
options flexscale.enable
```

2. To enable or disable the WAFL external cache, enter the following command:

```
options flexscale.enable {on|off}
```

Caching normal user data blocks

If you cache normal user data blocks, the WAFL external cache interprets this setting as the buffer cache policy of `keep` and saves normal user data blocks in the external 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 external cache interprets this setting as the buffer cache policy of `keep` and saves normal user data blocks in the external 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 255

Caching low-priority user data blocks

You can cache low-priority user data blocks that are not normally stored by WAFL external 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 the iSCSI, NFS, or CIFS protocols.

About this task

Caching low-priority user data blocks is useful if you have workloads that fit within WAFL external cache memory and if the workloads consist of either write followed by read or large sequential reads.

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 262

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 external 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 external cache interprets this setting as the buffer cache policy of `reuse` and does not save normal data blocks or low-priority blocks in the external cache.

Steps

1. Enter the following command to turn off normal user data block caching:

```
options flexscale.normal_data_blocks off
```
2. Enter the following command to turn off low-priority user data block caching:

```
options flexscale.lopri_blocks off
```

Related concepts

[FlexShare and the buffer cache policy values](#) on page 255

Related tasks

[Using FlexShare to set volume buffer cache policy](#) on page 258

Displaying the WAFL external cache configuration

Data ONTAP enables you to display configuration information for WAFL external cache.

Step

1. Enter the following command:

```
stats show -p flexscale
```

Displaying usage and access information for WAFL external cache

You can display usage and access information for WAFL external 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 external cache once every 10 seconds for 5 times:

```
stats show -p flexscale-access -i 10 -n 5
```

Preserving the cache in the Flash Cache family of modules

The system does not serve data from a Flash Cache or Flash Cache 2 module when a node is shutdown. However, the WAFL external cache preserves the cache during a graceful shutdown and can serve "warm" data after giveback.

The WAFL external cache can preserve the cache in Flash Cache modules during a graceful shutdown. It preserves the cache through a process called "cache rewarming," which helps to maintain system performance after a graceful shutdown. For example, you might shut down a system to add hardware or upgrade software.

Cache rewarming is enabled by default if you have a Flash Cache or Flash Cache 2 module installed. Cache rewarming is available when both nodes in an HA pair are running Data ONTAP 8.1 or later.

Related concepts

Increasing WAFL cache memory on page 260

How cache rewarming works

WAFL external cache initiates the cache rewarming process during a reboot or a takeover and giveback. The process keeps the cache in Flash Cache and Flash Cache 2 modules "warm."

When a storage system powers down, the WAFL external cache takes a snapshot of the data in Flash Cache and Flash Cache 2 modules. When the system powers up, it uses the snapshot to rebuild the cache. After the process completes, the system can read data from the cache.

In an HA configuration, cache rewarming is more successful when minimal changes are made to data during takeover and giveback. When you initiate takeover and giveback, the takeover partner maintains a log of data written to the down partner's storage. If there are changes to a large amount of the data that is stored in the cache, then the cache rewarming process has more data to rewarm when the node comes back online. As a result, the cache may require additional warming time.

Note: Cache rewarming does not work if the WAFL external cache functionality is disabled.

Events that initiate cache rewarming

You can initiate cache rewarming when you shut down a node or when you initiate takeover and giveback.

The following commands initiate cache rewarming:

- `halt [-t]`
- `halt [-t] -f`
- `reboot [-t]`
- `reboot [-t] -s`
- `reboot [-t] -f`
- `cf takeover [-f] [-n]`

Events that do not initiate cache rewarming

WAFL external cache does not initiate cache rewarming if the storage system crashes, if there is a sudden loss of power, or if you run certain commands.

The following commands do not initiate cache rewarming:

- `halt -d`
- `reboot -d`
- `cf forcetakeover [-f]`

Events that abort cache rewarming

After the cache rewarming process starts, some events can abort the entire process and some events can abort the process on specific aggregates.

The following events abort the entire cache rewarming process:

- You add, remove, or move a Flash Cache or Flash Cache 2 module after the WAFL external cache takes the snapshot, but before it rebuilds the cache.
- The takeover node crashes.
- The local node crashes as the WAFL external cache rebuilds the cache.
- After a node reboots, it shuts down before the WAFL external cache can rebuild the cache.
- You initiate a SnapRestore operation on the node's root aggregate before the WAFL external cache rebuilds the cache.
- The `wafliiron` process mounts the root aggregate.

The following events abort cache rewarming on the affected aggregate:

- You initiate a SnapRestore operation on an aggregate before the WAFL external cache rebuilds the cache.
- An aggregate does not come online within 20 minutes after the WAFL external cache starts to rebuild the cache.
- The `wafliiron` process mounts an aggregate.

Enabling and disabling cache rewarming

Cache "rewarming" is enabled by default if a Flash Cache or Flash Cache 2 module is installed. You can disable and then re-enable cache rewarming, if necessary. You should do this only under the guidance of technical support.

Before you begin

You can enable cache rewarming if the following is true:

- A Flash Cache or Flash Cache 2 module is installed.
- The WAFL external cache functionality is enabled.

About this task

Cache rewarming works at the node level. To ensure that cache rewarming works during a takeover and giveback, enable it on all nodes.

Step

1. Enter one of the following commands:

If you want to...	Use this command:
Disable cache rewarming	<code>options flexscale.rewarm off</code>
Enable cache rewarming	<code>options flexscale.rewarm on</code>

Related tasks

[Enabling and disabling WAFL external cache](#) on page 261

Optimizing 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.

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.

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 five 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 guarantee is enabled so that the reallocation scan does not cause an over-commitment of the volume's storage space.
For information about volume guarantees, see the *Data ONTAP Storage Management Guide for 7-Mode*.
- 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 269

[Managing reallocation scans](#) on page 270

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 example, 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: Do not run an aggregate reallocation scan if free space reallocation is enabled 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.

Related concepts

[How free space reallocation optimizes free space](#) on page 282

Related references

[Commands for managing free space reallocation](#) on page 284

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 qtrees 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.

Managing 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:

1. Enable reallocation scans.
2. Do one of the following:
 - a. Define a reallocation scan to run at specified intervals (such as every 24 hours).
 - b. 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 re-create its reallocation scan. If the file or LUN has not been re-created 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.

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] [-o] [-p] [-u] [-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.
- `-o` runs the job once and then automatically removes it from the system.
- `-p` reallocates user data on the physical blocks in the aggregate while preserving the logical block locations within a FlexVol volume. You cannot use the `-p` option with the `-u` option. This option also reallocates the shared blocks in a deduplicated volume. Reallocation scans skip deduplicated data if you do not specify the `-p` option. You can use this option only with FlexVol volumes, with files and LUNs within FlexVol volumes, or with deduplicated volumes.

Do not use `-p` when you start a reallocation scan on a compressed volume. Starting a reallocation scan on a compressed volume using `-p` does not optimize the layout of a volume.

- `-u` duplicates blocks that are shared between files by deduplication. Duplicating the blocks removes the sharing. This option can help remove fragmentation; however, because blocks are duplicated, it can result in increased disk usage, especially for full reallocation. You cannot use the `-u` option with the `-p` option.
- `-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

[Managing reallocation scans](#) on page 270

Related tasks

[Creating a reallocation scan schedule](#) on page 273

[Enabling reallocation scans](#) on page 271

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.

About this task

An aggregate reallocation scan reallocates free space and is not the same as 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.

Note: Do not run an aggregate reallocation scan if free space reallocation is enabled 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.

Step

1. On the command line for the storage system, 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 268

[How free space reallocation optimizes free space](#) on page 282

Related tasks

[Performing a full reallocation scan of a LUN, file, or volume](#) on page 275

[Creating a reallocation scan schedule](#) on page 273

Related references

[Commands for managing free space reallocation](#) on page 284

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.

About this task

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 p.m.

```
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.

About this task

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	reallocate start -o -n <i>pathname</i>
An aggregate	reallocate start -A -o <i>aggr_name</i>

- **-o** performs the scan only once.
- **-n** performs the scan without checking the layout of the LUN, file, or volume.

Example

The following command syntax 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.

About this task

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] [-u] 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 cannot use the `-p` option with the `-u` option. This option also reallocates the shared blocks in a deduplicated volume. Reallocation scans skip deduplicated data if you do not specify the `-p` option. You can use this option only with FlexVol volumes, with files and LUNs within FlexVol volumes, or with deduplicated volumes.
Do not use `-p` when you start a reallocation scan on a compressed volume. Starting a reallocation scan on a compressed volume using `-p` does not optimize the layout of a volume.
- `-u` duplicates blocks that are shared between files by deduplication. Duplicating the blocks removes the sharing. This option can help remove fragmentation; however, because blocks are duplicated, it can result in increased disk usage, especially for full reallocation. You cannot use the `-u` option with the `-p` option.

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.

About this task

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 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 command syntax 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
```

Result

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. 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.

About this task

When you restart a file, LUN, or volume reallocation scan, the scan restarts from the beginning of the reallocation process. An aggregate reallocation scan restarts from where it stopped.

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.

About this task

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.

Improving read performance

You can improve the read performance of your storage system by enabling read reallocation on volumes. Read reallocation is disabled by default.

What read reallocation is

For workloads that perform a mixture of random writes and large and multiple sequential reads, read reallocation improves file layout and sequential read performance.

Read reallocation analyzes the parts of the file that are read sequentially. If the associated blocks are not already largely contiguous, Data ONTAP updates the layout by rewriting those blocks to another location on disk. The rewrite improves the 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.

Read reallocation is not supported on compressed volumes and FlexCache volumes.

Enabling or disabling read reallocation

You can enable read reallocation on traditional volumes and FlexVol volumes to improve subsequent read performance of a file. You can also disable read reallocation.

About this task

If file fragmentation is a concern for FlexCache volumes, enable read reallocation on the FlexCache origin volume.

Step

1. Enter the following command:

```
vol options vol-name read_realloc [on | space_optimized | off]
```

- `on` enables read reallocation on the volume.
- `space_optimized` also enables read reallocation but can be used only on FlexVol volumes or deduplicated 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 reallocation 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 also result in degraded Snapshot read performance.

Use `space_optimized` with deduplicated volumes to rearrange the volume's shared blocks. Read reallocation skips deduplicated data if you set the `read_realloc` option to `on`.

`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 and is the default setting.

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

Example

The following command enables read reallocation on the `vol1` volume:

```
vol options vol1 read_realloc space_optimized
```

Improving write performance


You can enable free space reallocation on aggregates to improve write performance. Free space reallocation improves write performance by optimizing the free space within an aggregate. Free space reallocation is disabled by default.

How free space reallocation optimizes free space

Free space reallocation optimizes the free space in an aggregate immediately before Data ONTAP writes data to the blocks in that aggregate.

Before Data ONTAP writes data to a segment of blocks in an aggregate, free space reallocation evaluates the layout of those blocks. If the layout is not optimal, the free space reallocation function rearranges the blocks. Rearranging the blocks increases the amount of contiguous free space available in the aggregate, which improves the performance of Data ONTAP writes to those blocks.

The following graphic illustrates how free space reallocation optimizes the free space in a segment of blocks:



When to enable free space reallocation

Free space reallocation works best on workloads that perform a mixture of small random overwrites and sequential or random reads. You can expect additional CPU utilization when you enable free space reallocation. You should not enable free space reallocation if your storage system has sustained, high CPU utilization.

Note: You can use the `sysstat` command to monitor CPU utilization.

For best results, you should enable free space reallocation when you create a new aggregate. If you enable free space reallocation on an existing aggregate, there might be a period where Data ONTAP performs additional work to optimize free space. This additional work can temporarily impact system performance.

When to use free space reallocation with other reallocation features

You can use free space reallocation with the other reallocation features that you use to manage system performance: reallocation scans and read reallocation. However, you should not run an aggregate reallocation scan (`reallocate -A`) if free space reallocation is enabled on the aggregate.

When you enable free space reallocation, you should also consider enabling read reallocation. Free space reallocation and read reallocation are complimentary technologies that optimize data layout. Read reallocation optimizes the system for sequential reads, while free space reallocation optimizes for writes.

Related concepts

[Optimizing LUN, file, volume, and aggregate layout](#) on page 267

[What read reallocation is](#) on page 281

How free space reallocation differs from an aggregate reallocation scan

Like an aggregate reallocation scan (`reallocate -A`), free space reallocation optimizes the location of blocks in an aggregate. However, an aggregate reallocation scan is a one-time scan of every segment within an aggregate, while free space reallocation runs continuously and scans only the segments to be used.

Types of aggregates that free space reallocation can and cannot optimize

Free space reallocation optimizes the free space in specific types of aggregates.

Free space reallocation optimizes free space in the following:

- Aggregates that provide storage to FlexVol volumes
- The HDD RAID groups in an aggregate

Free space reallocation does not optimize free space in the following:

- Aggregates that provide storage to traditional volumes
- The SSD RAID groups in an aggregate
- Aggregate Snapshot copies
- Read-only volumes

Commands for managing free space reallocation

Use the `aggr options` command to manage free space reallocation.

If you want to...	Use this command...
Enable free space reallocation on an aggregate	<code>aggr options aggr_name free_space_realloc on</code>

If you want to...	Use this command...
Disable free space reallocation on an aggregate	<code>aggr options <i>aggr_name</i> free_space_realloc off</code>
Identify whether free space reallocation is enabled or disabled on an aggregate	<code>aggr options <i>aggr_name</i></code>

For more information, see the man pages.

Troubleshooting tools

If you experience problems with your storage system, some tools are available to help you understand and avoid problems.

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 NetApp Support 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

NetApp Support Site: support.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.

Related tasks

[Configuring message logging](#) on page 144

Using the Syslog Translator to get 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 NetApp Support 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

[NetApp Support Site: support.netapp.com](http://support.netapp.com)

How to use the NetApp Support Site for help with errors

The NetApp Support Site is a powerful resource to help you diagnose and solve problems with your storage system.

The NetApp Support 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

[NetApp Support Site: support.netapp.com](http://support.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

[Troubleshooting the storage system by using the RLM](#) on page 231

Related references

[Troubleshooting a system by using the SP](#) on page 206

Glossary

A

ACL	Access control list.
active/active configuration	<ul style="list-style-type: none"> In the Data ONTAP 7.2 and 7.3 release families, a pair of storage systems or V-Series systems (sometimes called <i>nodes</i>) configured to serve data for each other if one of the two systems stops functioning. Also sometimes referred to as <i>active/active pairs</i>. In the Data ONTAP 8.x release family, this functionality is referred to as a <i>high-availability (HA) configuration</i> or an <i>HA pair</i>. In the Data ONTAP 7.1 release family and earlier releases, this functionality is referred to as a <i>cluster</i>.
address resolution	The procedure for determining an address corresponding to the address of a LAN or WAN destination.
admin Vserver	In Data ONTAP operating in Cluster-Mode, a Vserver that has overall administrative access to all objects in the cluster, including all objects owned by other Vservers, but does not provide data access to clients or hosts.
administration host	A client computer that is used to manage a storage system through a Telnet or Remote Shell connection.
Application Program Interface (API)	A language and message format used by an application program to communicate with the operating system or some other system, control program, or communications protocol.
authentication	The process of verifying the identity of a user who is logging in to a computer system.
AutoSupport	An integrated technology that triggers email messages from the customer site to technical support or another specified email recipient when there are any failures in Unified Manager services. These messages contain information such as feature usage metrics, configuration and user settings, system health, and so on.

B

big-endian	A binary data format for storage and transmission in which the most significant byte comes first.
-------------------	---

C

caching module	A Flash Cache 2, Flash Cache, or Performance Acceleration Module (PAM) PCIe-based, memory module that optimizes the performance of random
-----------------------	---

read-intensive workloads by functioning as an intelligent external read cache. This hardware works in tandem with the WAFL External Cache software component of Data ONTAP.

CIFS share	<ul style="list-style-type: none"> • In Data ONTAP, a directory or directory structure that has been made available to network users and can be mapped to a drive letter on a CIFS client. Also known simply as a <i>share</i>. • In OnCommand Insight (formerly SANscreen suite), a service exposed from a NAS device to provide file-based storage through the CIFS protocol. CIFS is mostly used for Microsoft Windows clients, but many other operating systems can access CIFS shares as well.
CLI	command-line interface. The storage system prompt is an example of a command-line interface.
client	A workstation or PC in a client-server architecture; that is, a computer system or process that requests services from and accepts the responses of another computer system or process.
cluster	<ul style="list-style-type: none"> • In clustered Data ONTAP 8.x, a group of connected nodes (storage systems) that share a namespace and that you can manage as a single virtual server or multiple virtual servers, providing performance, reliability, and scalability benefits. • In the Data ONTAP 7.1 release family and earlier releases, a pair of storage systems (sometimes called <i>nodes</i>) configured to serve data for each other if one of the two systems stops functioning. • In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an <i>active/active configuration</i>. • For some storage array vendors, <i>cluster</i> refers to the hardware component on which host adapters and ports are located. Some storage array vendors refer to this component as a <i>controller</i>.
cluster Vserver	Previous name for a <i>data Vserver</i> . See <i>data Vserver</i> .
Common Internet File System (CIFS)	Microsoft's file-sharing networking protocol that evolved from SMB.
community	A logical relationship between an SNMP agent and one or more SNMP managers. A community is identified by name, and all members of the community have the same access privileges.
console	The physical or virtual terminal that is used to monitor and control a storage system.
Copy-On-Write (COW)	The technique for creating Snapshot copies without consuming excess disk space.
D	

data Vserver	In clustered Data ONTAP, a virtual server that facilitates data access from the cluster; the hardware and storage resources of the cluster are dynamically shared by data Vservers within a cluster. Previously referred to as a <i>cluster Vserver</i> .
degraded mode	The operating mode of a storage system when a disk in the RAID group fails or the batteries on the NVRAM card are low.
disk ID number	The number assigned by the storage system to each disk when it probes the disks at startup.
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.
E	
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. Sometimes called an <i>adapter</i> .
expansion slot	The slots on the storage system board into which you insert expansion cards.
F	
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.
Flash Cache/Flash Cache 2	The two members of the family of PCIe-based, solid state memory modules that optimize the performance of random read-intensive workloads by functioning as an intelligent external read cache. The Flash Cache 2 module is the successor of the Flash Cache module, which is the successor of the Performance Acceleration Module (PAM). This hardware works in tandem with the WAFL External Cache software component of Data ONTAP.
G	
giveback	The technology that enables two storage systems to return control of each other's data after the issues that caused a controller failover are resolved.
global namespace	See <i>namespace</i> .
group	In Data ONTAP operating in 7-Mode, a group of users defined in the storage system's <code>/etc/group</code> file.

Group ID (GID)	The number used by UNIX systems to identify groups.
H	
HA (high availability)	<ul style="list-style-type: none"> • In Data ONTAP 8.x, the recovery capability provided by a pair of nodes (storage systems), called an <i>HA pair</i>, that are configured to serve data for each other if one of the two nodes stops functioning. • In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an <i>active/active configuration</i>.
HA pair	<ul style="list-style-type: none"> • In Data ONTAP 8.x, a pair of nodes whose controllers are configured to serve data for each other if one of the two nodes stops functioning. Depending on the system model, both controllers can be in a single chassis, or one controller can be in one chassis and the other controller can be in a separate chassis. • In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an <i>active/active configuration</i>.
heartbeat	A repeating signal transmitted from one storage system to the other that indicates that the storage system is in operation. Heartbeat information is also stored on disk.
hot 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.
I	
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.
L	
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.
M	
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	In Data ONTAP operating in 7-Mode, 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.

N

namespace	In network-attached storage (NAS) environments, a collection of files and path names to the files.
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 card.
node Vserver	In Data ONTAP operating in Cluster-Mode, a Vserver that is restricted to operation in a single node of the cluster at any one time, and provides administrative access to some objects owned by that node. A node Vserver does not provide data access to clients or hosts.
normal mode	The state of a storage system when there is no takeover in the high-availability configuration.
NVMEM	nonvolatile memory.
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 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.

P

PAM (Performance Acceleration Module)	A PCIe-based, DRAM memory module that optimizes the performance of random read-intensive workloads by functioning as an intelligent external read cache. This hardware is the predecessor of the Flash Cache module and works in tandem with the WAFL External Cache software component of Data ONTAP.
panic	A serious error condition causing the storage system or V-Series 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 mode	The method you use to communicate through the command-line interface with a virtual storage system during a takeover.
partner node	From the point of view of the local node (storage system), the other node in a high-availability configuration.

Performance Acceleration Module (PAM)	See <i>PAM (Performance Acceleration Module)</i> .
POST	Power-on self-tests. The tests run by a storage system after the power is turned on.
Q	
qtree	A special subdirectory of the root of a volume that acts as a virtual subvolume with special attributes.
R	
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 RAID4, 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.
S	
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.
SFO	See <i>storage failover (SFO)</i> .
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.
storage failover (SFO)	In Data ONTAP operating in Cluster-Mode, the method of ensuring data availability by transferring the data service of a failed node to another node in an HA pair. Transfer of data service is often transparent to users and applications. In Data ONTAP 7.2 and later, and in Data ONTAP operating in 7-Mode, the failover method is called <i>controller failover</i> .

T

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 node (storage system) when it has taken over its partner. The console prompt indicates when the node is in takeover mode.
takeover node	A node (storage system) that remains in operation after the other node stops working and that hosts a virtual node 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.

U

UID	user identification number.
Unicode	A 16-bit character set standard. It was designed and is maintained by the nonprofit consortium Unicode Inc.

V

vFiler unit	In Data ONTAP operating in 7-Mode, a virtual storage system that 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.
Vserver	In Data ONTAP operating in Cluster-Mode, a virtual server that provides network access through unique network addresses, that might serve data out of a distinct namespace, and that is separately administrable from the rest of the cluster. There are three types of Vservers— <i>admin</i> , <i>data</i> , and <i>node</i> —but unless there is a specific need to identify the type of Vserver, <i>Vserver</i> usually refers to <i>data</i> Vserver.

W

WAFL	Write Anywhere File Layout. A file system designed for the storage system to optimize write performance.
WAFL External Cache	On a storage system that has a Performance Acceleration Module (PAM), Flash Cache, or Flash Cache 2 module installed, this cache improves storage system performance by reducing the number of disk reads. Sometimes referred to as <i>WAFL extended cache</i> .
WINS	Windows Internet Name Service.

workgroup	A collection of computers running Windows NT or Windows for Workgroups that is grouped for browsing and sharing.
------------------	--

Copyright information

Copyright © 1994–2013 NetApp, Inc. All rights reserved. Printed in the U.S.

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. 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, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, FAServer, FastStak, FilerView, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

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

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

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

NetApp, Inc. NetCache is certified RealSystem compatible.

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 email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

- /etc directory
 - access with FTP [79](#)
 - access with SFTP [80](#)
- /etc/log/auditlog file [146](#)
- /etc/messages file [78](#)
- /etc/rc
 - recovering from errors [149](#)
- /etc/rc file
 - commands in [146](#)
 - sample [147](#)
- /etc/syslog.conf file
 - configuring message logging in [144](#)
 - file format and parameters of [142](#)
- /etc/usermap.cfg file, character coding of [78](#)
- /home directory
 - access with FTP [81](#)
 - access with SFTP [82](#)
- /home file, contents of [74](#)

A

- access
 - allowing only secure to storage systems [69](#)
 - disabling root account for the storage system [98](#)
 - restricting protocol [65](#)
 - restricting RLM to only the specified administration hosts [223](#)
 - restricting SP to only the specified administration hosts [209](#)
 - secure protocols and storage system [35](#)
- accessing
 - /etc directory with FTP [79](#)
 - /etc directory with SFTP [80](#)
 - /home directory with FTP [81](#)
 - /home directory with SFTP [82](#)
 - data management, introduction [17](#)
 - starting a Telnet session [53](#)
 - storage system from Windows client with RSH [58](#)
 - storage system with serial port [33](#)
 - the SP from an administration host [194](#)
 - the system with the remote management device [34](#)
- accounts
 - enabling or disabling diagnostic [128](#)
 - reasons for creating administrator [95](#)
- administering
 - systems, methods for [28](#)
- administration
 - methods for storage systems [22](#)
- administration hosts
 - accessing the SP from [194](#)
 - adding [64](#)
 - defined [62](#)
 - how to specify [63](#)
 - reasons for designating workstations as [62](#)
 - removing [64](#)
 - restricting RLM access to only the specified [223](#)
 - restricting SP access to only the specified [209](#)
 - use of [62](#)
- administrative level commands [25](#)
- administrator accounts
 - changing the password of (passwd) [124](#)
 - reasons for creating [95](#)
- aggregate information
 - displaying [238](#)
- aggregate Snapshot copies
 - considerations for managing [132](#)
- aggregate Snapshot reserve
 - considerations for increasing [133](#)
- aggregates
 - improving write performance [282](#)
 - introduction to root [71](#)
 - root option [83](#)
- alerts
 - .See health monitoring
- attachment
 - AutoSupport messages [162](#)
- audit logging
 - introduction [145](#)
 - logs
 - audit, introduction [145](#)
- authentication
 - public-key-based overview [41](#)
 - with SSL [46](#)
- AutoSupport
 - about [157](#)
 - commands [168](#)
 - communication with technical support [160](#)
 - configuring [168](#)
 - content [161–163](#)
 - daily messages [157, 161](#)
 - defined [157](#)

- displaying information [168](#)
- email [163](#)
- enabled by default
 - having messages sent to your organization [157](#)
- enabling and disabling [168](#)
- event-triggered [162](#)
- event-triggered messages [157](#), [161](#)
- events [162](#)
- files [163](#)
- getting message descriptions [168](#)
- history [168](#)
- log files [162](#)
- mail host support for [164](#)
- manifest [168](#)
- Message Matrices [168](#)
- modifying triggers [168](#)
- options [169](#)
- performance messages [157](#), [161](#)
- requirements for [164](#)
- resending messages [168](#)
- sending messages [168](#)
- setup [166](#)
- severity types [164](#)
- subsystems [161](#), [162](#)
- transport protocol [164](#)
- troubleshooting
 - HTTP [178](#)
 - HTTPS [178](#)
 - SMTP [179](#)
- troubleshooting mail host relaying [180](#)
- troubleshooting messages [177](#)
- weekly messages [157](#), [161](#), [163](#)
- when messages are generated [157](#)
- when sent [157](#)
- where sent [157](#)

AutoSupport manifest

- content of [176](#)

B

- banner messages
 - configuration [54](#)
- block sharing
 - multiprotocol [16](#)
- boot devices
 - recovering from a corrupted image of the system's [88](#)
 - storage system [86](#)
- boot environment prompt
 - booting Data ONTAP from [87](#)

- boot menus
 - managing the storage system with the [92](#)
- boot options
 - storage system [86](#)
- booting
 - Data ONTAP at the boot environment prompt [87](#)
 - storage systems [86](#)
- browsers, improving security through [46](#)
- buffer cache policy
 - setting [258](#)

C

- CA
 - installing certificates signed by [48](#)
- CA signed certificates
 - installing [48](#)
- cache rewarming
 - abort events [266](#)
 - about [264](#)
 - disabling [266](#)
 - enabling [266](#)
 - how it works [265](#)
 - trigger events [265](#)
- caches
 - comparison of Flash Pool and Flash Cache [260](#)
- callhome events [162](#)
- capabilities
 - definition of [95](#)
 - example of creating a user with custom [118](#)
 - how users are assigned [96](#)
 - list of supported types [108](#)
 - modifying others' [102](#)
 - of predefined roles [106](#)
 - types of [108](#)
- Certificate Authority
 - See* CA
- certificates
 - domain names and [49](#)
 - generating [47](#)
 - installing signed CA [48](#)
- change privileges, file ownership [67](#)
- character coding for configuration files [78](#)
- CIFS
 - accessing /etc directory [79](#)
 - accessing /home directory [81](#)
 - client, requirements to manage storage system [63](#)
 - editing configuration files using [77](#)
- CIFS share lookups
 - controlling anonymous [67](#)

CLI

- rules for using Data ONTAP [23](#)

clients

- accessing storage system from Windows, with RSH [58](#)
- CIFS, requirements [63](#)
- editing configuration file from [77](#)
- NFS, requirements [63](#)

command-line editor

- using the [24](#)

command-line interface

- See* CLI

commands

- administrative level [25](#)
- advanced level [25](#)
- date (sets system date and time) [136](#)
- displaying history [24](#)
- for managing a system at the SP admin privilege level [198](#)
- for managing a system at the SP advanced privilege level [201](#)
- for managing licenses [135](#)
- for managing the RLM [232](#)
- for managing the SP with Data ONTAP [210](#)
- halt (halts the storage system) [91](#)
- options waf1.root_only_chown (sets file ownerships changes) [67](#)
- passwd (changes administrative user password) [124](#)
- passwd (changes storage system system password) [123](#)
- privilege level [25](#)
- privilege levels [25](#)
- reboot (reboots the storage system) [90](#)
- RSH command list [59](#)
- stats [242](#)
- timezone (displays and sets system time zone) [138](#)

CompactFlash cards

- checking the Data ONTAP version of [89](#)

configuration files

- /etc [74](#)
- backing up [150](#)
- backing up and cloning [149](#)
- cloning [151](#)
- comparing backups [152](#)
- editing from CIFS client [77](#)
- editing from NFS client-setup [77](#)
- hard limits [76](#)
- introduction [75](#)
- restoring [151](#)
- within /etc directory [75](#)

configuration information

- displaying storage system [236](#)

configurations

- of banner messages [54](#)

considerations

- for increasing the aggregate Snapshot reserve [133](#)
- for managing core dump files [139](#)

console sessions

- RLM CLI and system [227](#)
- SP CLI and system [197](#)

consoles

- accessing the serial console from the RLM [226](#)
- accessing the serial console from the SP [196](#)
- accessing the SP from serial [195](#)
- session rules for Telnet, SSH-interactive, and [29](#)

core dump files

- considerations for managing [139](#)
- managing [139](#)
- methods of segmenting [139](#)

core segments

- commands for managing [140](#)

custom capabilities

- example of creating a user with [118](#)

D

- daily AutoSupport messages [157](#), [161](#)

data

- Data ONTAP functionality for managing migration [17](#)

data access

- management, introduction [17](#)

Data ONTAP

- booting at the boot environment prompt [87](#)
- checking version of [89](#)
- management tasks by using e0M [32](#)
- restoring LUNs [132](#)
- rules for using CLI [23](#)

Data ONTAP commands

- for managing the SP [210](#)

Data ONTAP-v

- components [15](#)

Data ONTAP-v systems

- root volume, introduction to [71](#)

data organization

- management [17](#)

data protection

- capabilities, defined [18](#)

data storage management [16](#)

- date, setting storage system time and [136](#)

- default directories [74](#)
- diagnostic accounts
 - enabling or disabling [128](#)
 - setting the password for [129](#)
 - uses of [128](#)
- directories, default permissions [74](#)
- disk shelves introduction to [14](#)
- DNS [178](#)
- domain names, changing storage system [49](#)
- domain users
 - definition of [95](#)
 - deleting [117](#)
 - granting access to [101](#)
 - listing [112](#)

E

- e0M interface
 - Data ONTAP management tasks by using [32](#)
 - introduction to [31](#)
- e0M interfaces
 - differences between the remote management device and [32](#)
- editor
 - using the command-line [24](#)
- email
 - AutoSupport [163](#)
- EMS
 - callhome event [159](#)
 - data in AutoSupport messages [162](#)
 - event-triggered AutoSupport messages, and [159](#)
 - managing event messages [155](#)
 - unknown user event [180](#)
- encryption
 - with SSL [46](#)
- environmental status
 - commands for displaying [240](#)
- error message logging, about [287](#)
- error messages
 - getting information by using Syslog Translator [287](#)
- errors
 - how to use NetApp Support Site for help with [288](#)
 - recovering from /etc/rc [149](#)
- event management
 - displaying EMS information [155](#)
 - displaying event log information [156](#)
- Event Management Systems
 - See* EMS
- event messages
 - managing [155](#)

- event-triggered AutoSupport messages
 - EMS, and [159](#)
 - files collected for message [176](#)
 - subsystems [161](#)
- events
 - AutoSupport messages [162](#)

F

- FAS systems
 - root volumes, introduction to [71](#)
- Fibre Channel
 - getting information [241](#)
- file ownership change privileges [67](#)
- file services
 - through network [15](#)
- file sharing
 - multiprotocol [16](#)
- File Transfer Protocol
 - See* FTP
- files
 - for configuration, introduction [75](#)
 - methods of segmenting core dump [139](#)
- files, configuration [74](#)
- firmware updates
 - methods of managing SP [207](#)
- Flash Cache
 - compared with Flash Pools [260](#)
- Flash Cache family of modules [260](#)
- Flash Pools
 - compared with Flash Cache [260](#)
- flexscale.rewarm option [266](#)
- FlexShare
 - about [252](#)
 - assigning system operation priorities [257](#)
 - assigning volume priorities [256](#)
 - buffer cache policy [255](#), [258](#)
 - default queue [255](#)
 - high-availability configuration [254](#)
 - io_concurrency options [255](#)
 - modifying default priority [259](#)
 - priority levels [254](#)
 - removing priorities [259](#)
 - volume operations and [254](#)
 - volume prioritization [255](#)
 - when to use [253](#)
- FlexVol volumes
 - sizing considerations for root [73](#)
- free space reallocation
 - disabling [284](#)

- enabling [284](#)
- how it works [282](#)
- overview [282](#)
- relation to aggregation reallocation scans [284](#)
- supported aggregates [284](#)
- using with other reallocation features [284](#)
- viewing status [284](#)
- when to enable [283](#)

FTP

- for accessing the /etc directory [79](#)
- for accessing the /home directory [81](#)

G

generating certificates [47](#)

groups

- assigning roles to [104](#)
- creating users and assigning to [99](#)
- definition of [95](#)
- deleting [117](#)
- granting access and mapping roles for LDAP [119](#)
- listing [112](#)
- naming requirements [96](#)
- predefined [103](#)
- reloading from lclgroups.cfg file [105](#)
- renaming [105](#)
- setting maximum auxiliary [106](#)
- Windows special [97](#)

H

hard limits, configuration files [76](#)

health monitoring

- commands [186](#)
- example of responding to degraded health [184](#)
- how alerts trigger AutoSupport messages and events [183](#)
- how it works [181](#)
- responding to degraded health [183](#)
- ways to control when alerts occur [182](#)
- ways to respond to alerts [181](#)
- what health monitors are available [183](#)
- what it is [181](#)

health monitors [155](#)

history

- displaying command [24](#)

hosts

- accessing the SP from administration [194](#)
- definition of [62](#)
- how to specify administration [63](#)

- restricting RLM access to only the specified administration [223](#)
- restricting SP access to only the specified administration [209](#)

HTTP access to log files [82](#)

I

images

- recovering from the corruption of the system's boot device [88](#)

increasing cache memory [260](#)

interface

- eOM, introduction [31](#)

K

key pairs

- generating for SSH 2.0 [42](#)
- generating RSA for SSH 1.x [41](#)

keys

- public-based, authentication overview [41](#)

L

lclgroups.cfg file, reloading [105](#)

LDAP groups

- granting access and mapping roles for [119](#)

license

- types [134](#)

licenses

- commands for managing [135](#)
- managing [134](#)

log files

- AutoSupport messages [162](#)

log files, accessing using HTTP or HTTPS [82](#)

logs

- introduction to message [142](#)

lookups

- controlling anonymous CIFS share [67](#)

LUN (Logical Unit Number)

- restore [132](#)

LUNs

- reallocating to improve performance [268](#)

M

mail host support for AutoSupport [164](#)

managing

- licenses [134](#)
- manifest
 - event-triggered AutoSupport messages, for [176](#)
- message files, accessing using HTTP or HTTPS [82](#)
- message logging
 - introduction [142](#)
- messages
 - managing event [155](#)
- methods
 - for administering systems [28](#)
- migration
 - Data ONTAP functionality for managing data [17](#)
- monitoring
 - node connectivity [181](#)
- mounts
 - controlling privilege of NFS [66](#)
- multiprotocols
 - for file and block sharing [16](#)

N

- naming requirements for useradmin command [96](#)
- NetApp Support Site
 - how to use for help with errors [288](#)
- network
 - configuring the SP [191](#)
 - file service [15](#)
- NFS
 - access to /etc directory [79](#)
 - access to /home directory [81](#)
- NFS client
 - requirements to manage storage system [63](#)
- NFS mounts
 - controlling privilege of [66](#)
- node connectivity health monitor
 - commands for [186](#)
 - what it is [183](#)
- nonlocal users, granting access to [101](#)
- notifications
 - automatic, technical support upon system reboots [141](#)
- NVRAM
 - halt command to save data to disk [91](#)

O

- obsolete domain names, and SSL [49](#)
- online command-line help [25](#)
- online help
 - for using RLM CLI [227](#)

- for using SP CLI [197](#)
- options
 - Data ONTAP password management [124](#)
 - of Data ONTAP for managing system time [137](#)
 - security [68](#)
- overview [281](#)
- ownership change privileges, file [67](#)

P

- PAM (Performance Acceleration Module) [260](#)
- panics
 - considerations for managing core dump files [139](#)
 - reacting to storage system [286](#)
- passwords
 - changing (passwd) [123](#)
 - Data ONTAP options for managing rules [124](#)
 - how to manage for security [121](#)
 - setting for the diagnostic account [129](#)
- perfmom, using to monitor performance [251](#)
- performance
 - improving write performance [282](#)
 - monitoring with perfmom [251](#)
 - read [281](#)
 - read reallocation [281](#)
- Performance Acceleration Module [260](#)
- performance AutoSupport messages [157](#), [161](#)
- performance improvements, in storage systems
 - WAFL external cache [260](#)
- perfstat
 - for getting system information [251](#)
- permissions of default directories (/etc, /home) [74](#)
- privilege levels for Data ONTAP commands [25](#)
- privileges
 - controlling NFS mount [66](#)
- privileges, file ownership change [67](#)
- prompts
 - booting Data ONTAP at the boot environment [87](#)
 - rebooting the storage system at the system [87](#)
- protocol
 - restricting access [65](#)
- protocols
 - for file and block sharing [16](#)
 - introduction to SSH [36](#)
 - secure, and storage system access [35](#)
- public-key-based
 - authentication overview [41](#)

Q

quota file, character coding for [78](#)

R

read reallocation

disabling [281](#)

enabling [281](#)

reading files [153](#)

reallocate commands

reallocate off [280](#)

reallocate on [271](#)

reallocate quiesce [278](#)

reallocate restart [278](#)

reallocate schedule [273](#)

reallocate start [271](#), [275](#)

reallocate start -A [272](#), [275](#)

reallocate status [279](#)

reallocate stop [279](#)

reallocate commands: reallocate schedule -d [274](#)

reallocation

best practices [280](#)

defining scans

aggregates [272](#)

LUNs, files, or volumes [271](#)

deleting a scan [279](#)

deleting scan schedule [274](#)

disabling scans [280](#)

enabling scans [271](#)

free space [282](#)

full [275](#)

managing scans [270](#)

measure-only [276](#)

quiescing scans [278](#)

read [281](#)

restarting scans [278](#)

scans [267](#)

scheduling scans [273](#)

starting one-time scan [275](#)

viewing scan status [279](#)

when to use with free space reallocation [284](#)

with LUNs, files, or volumes [268](#)

rebooting

the storage system at the system prompt [87](#)

the storage system remotely [90](#)

rebooting the system

from the console [90](#)

recovering

from a corrupted image of the system's boot device

[88](#)

reinitialization

of SSL [49](#)

remote

system management by using the RLM [215](#)

system management by using the SP [189](#)

Remote LAN Modules

See RLM

remote management

of a storage system [189](#)

remote management devices

accessing the system with [34](#)

differences between the e0M interface and [32](#)

Remote Shell (RSH) [56](#)

remote shell application

See see RSH

remotely

rebooting the storage system [90](#)

repeat mode

using stats command interactively in [246](#)

requests

issuing SSH [44](#)

requirements

to manage storage system on NFS clients [63](#)

reserves

considerations for increasing aggregate Snapshot

[133](#)

RLM

down filer events [233](#)

down system events [233](#)

logging in to [221](#)

managing with Data ONTAP commands [232](#)

SNMP traps [233](#), [234](#)

troubleshooting connection problems [235](#)

RLMs

accessing the serial console from [226](#)

CLI and system console sessions [227](#)

commands for managing [232](#)

commands for managing at the admin privilege level

[228](#)

commands for managing at the advanced privilege level [230](#)

commands for troubleshooting the storage system [231](#)

configuring [218](#)

configuring automatic logout of idle SSH

connections to [224](#)

Data ONTAP

- enabling or disabling SNMP traps for RLM and [234](#)
 - enabling or disabling SNMP traps for Data ONTAP and [234](#)
 - introduction to [216](#)
 - managing a system remotely by using [215](#)
 - remote management
 - RLMs [216](#)
 - restricting access to only the specified administration hosts [223](#)
 - using online help at CLI [227](#)
- RLMsways to configure [217](#)
- roles
 - assigning to groups [104](#)
 - creating [110](#)
 - definition of [95](#)
 - deleting [117](#)
 - listing [112](#)
 - mapping to LDAP groups [119](#)
 - modifying [111](#)
 - naming requirements [96](#)
 - predefined by Data ONTAP [106](#)
- root
 - disabling account access to the storage system [98](#)
- root aggregates
 - introduction to [71](#)
- root FlexVol volumes
 - sizing considerations for [73](#)
- root option for aggregates [83](#)
- root password, changing [123](#)
- root volume
 - changing [83](#)
 - directories contained within [74](#)
- root volumes
 - introduction to [71](#)
 - recommendations for [72](#)
- RSA
 - generating key pairs for SSH 1.x [41](#)
- RSH
 - accessing storage system from Windows client using [58](#)
- RSH (Remote Shell)
 - access to storage system [56](#)
- RSH commands
 - accessing storage system from a UNIX client [57](#)
 - displaying session information [60](#)
 - list of [59](#)
 - privilege levels [26](#)
 - use with user names and passwords [56](#)
- rules

- Data ONTAP options for managing password [124](#)

S

- scans, reallocation [270](#)
- secure access
 - allowing only to storage systems [69](#)
- secure protocols
 - and storage system access [35](#)
- Secure Shell
 - See* SSH
- Secure Sockets Layer
 - See* SSL
- SecureAdmin
 - displaying status of [52](#)
 - improving security with SSL [46](#)
 - managing SSH portion [38](#)
- security
 - controlling file ownership changes (options wafl.root_only_chown) [67](#)
 - how to manage passwords for [121](#)
 - limiting Telnet access [65](#)
 - settings [35](#)
- serial consoles
 - accessing from the RLM [226](#)
 - accessing from the SP [196](#)
 - accessing the SP from [195](#)
- serial ports
 - and slots of a storage system [13](#)
 - using to access storage system [33](#)
- Service Processors
 - See* SP
- sessions
 - terminating Telnet [54](#)
- setup
 - AutoSupport [166](#)
- severity
 - AutoSupport [164](#)
- SFTP
 - for accessing the /etc directory [80](#)
 - for accessing the /home directory [82](#)
- shares
 - controlling anonymous CIFS lookup [67](#)
- shelves
 - for disks [14](#)
- singleton mode
 - using stats command interactively in [245](#)
- sizing considerations
 - for root FlexVol volumes [73](#)
- slots

- and serial ports of a storage system [13](#)
- SMTP [179](#)
- Snapshot copies
 - considerations for managing aggregate [132](#)
- Snapshot reserve
 - considerations for increasing aggregate [133](#)
- SNMP traps
 - enabling or disabling for Data ONTAP and RLM [234](#)
- SP
 - logging in to [193](#)
 - sensors, discrete [204](#)
 - sensors, threshold-based [202](#)
 - SNMP traps [214](#)
- special system files
 - .bplustvoc_internal [132](#)
 - .vtoc_internal [132](#)
- SPs
 - accessing from an administration host [194](#)
 - accessing from the serial console [195](#)
 - accessing the serial console from [196](#)
 - CLI and system console sessions [197](#)
 - commands for managing a system at the admin privilege level [198](#)
 - commands for managing a system at the advanced privilege level [201](#)
 - commands for troubleshooting the storage system [206](#)
 - configuring automatic logout of idle SSH connections to [210](#)
 - configuring the network [191](#)
 - Data ONTAP commands for managing [210](#)
 - managing a system remotely by using [189](#)
 - methods of managing firmware updates [207](#)
 - restricting access to only the specified administration hosts [209](#)
 - using online help at CLI [197](#)
- SSH
 - configuring automatic logout of idle connections to the RLM [224](#)
 - configuring automatic logout of idle connections to the SP [210](#)
 - enabling or disabling [41](#)
 - generating key pairs for version 2.0 [42](#)
 - generating RSA key pairs for 1.x [41](#)
 - reinitializing [40](#)
 - setting up and starting [39](#)
- SSH (Secure Shell) commands
 - secureadmin status [52](#)
- SSH (Secure Shell) protocol
 - managing [38](#)
- SSH File Transfer Protocol
 - See* SFTP
- SSH interactive
 - configuring a timeout period [55](#)
 - controlling the timeout period [55](#)
- SSH protocol
 - introduction to [36](#)
- SSH requests
 - issuing [44](#)
- SSH-interactive
 - session rules for console, Telnet, and [29](#)
- SSL
 - enabling or disabling [50](#)
 - how to manage [47](#)
- SSL (Secure Sockets Layer) commands
 - secureadmin disable all [52](#)
 - secureadmin disable ssl [49](#)
 - secureadmin enable all [52](#)
 - secureadmin enable ssl [49](#)
 - secureadmin setup ssl [47](#)
 - secureadmin status [52](#)
- SSL (Secure Sockets Layer) protocol
 - authentication with [46](#)
 - enabling or disabling [49](#)
 - improving security with [46](#)
 - reinitializing [49](#)
 - setting up and starting [47](#)
- statistics commands
 - displaying Fibre Channel driver statistics [241](#)
 - displaying link statistics [241](#)
 - displaying relative physical drive position [241](#)
 - SAS statistics, description of [241](#)
- stats command
 - about [242](#)
 - background mode [247](#)
 - controlling output [248](#)
 - counters [242](#)
 - instances [242](#)
 - objects [242](#)
 - preset files [250](#)
 - using interactively in repeat mode [246](#)
 - using interactively in singleton mode [245](#)
- status
 - commands for displaying environmental [240](#)
- status commands
 - sasadmin (displays SAS adapter and expander information) [241](#)
 - sasstat adapter_state (displays state of a logical adapter) [241](#)

- sasstat dev_stats (displays statistics for disk drives connected to SAS channels) [241](#)
 - sasstat expander (displays SAS expander configuration) [241](#)
 - sasstat expander_map (displays SAS expander product information) [241](#)
 - sasstat expander_phy_state (displays SAS expander physical state) [241](#)
 - sasstat shelf (displays pictorial representation of the drive population of a shelf) [241](#)
 - sasstat shelf_short (displays the short form of the sasstat shelf command output) [241](#)
- status, displaying SecureAdmin [52](#)
- storage system
 - access and secure protocols [35](#)
- storage system access
 - /etc directory, accessing by using CIFS [79](#)
 - /etc directory, accessing by using NFS [79](#)
 - /home directory, accessing by using CIFS [81](#)
 - /home directory, accessing by using NFS [81](#)
 - using RSH from a UNIX client [57](#)
- storage systems
 - accessing from Windows client with RSH [58](#)
 - allowing only secure access to [69](#)
 - changing domain name of [49](#)
 - components of [12](#)
 - controlling file ownership changes (options wafl.root_only_chown) [67](#)
 - defined [12](#)
 - disabling root account access to [98](#)
 - displaying configuration information [236](#)
 - editing boot configuration file in [148](#)
 - halting (halt) [91](#)
 - how to boot [86](#)
 - internal components of, defined [12](#)
 - limiting Telnet access [65](#)
 - managing with the boot menu [92](#)
 - methods for administering [22](#)
 - reacting to panics [286](#)
 - rebooting at the system prompt [87](#)
 - rebooting remotely [88](#)
 - rebooting the system (reboot) [90](#)
 - remotely managing [189](#)
 - remotely rebooting [90](#)
 - RSH (Remote Shell) access to [56](#)
 - security [68](#)
 - setting date and time (date) [136](#), [138](#)
 - slots and serial ports of [13](#)
 - systems
 - remotely managing storage [189](#)
 - using serial port to access [33](#)
 - ways to boot [86](#)
- subsystems
 - AutoSupport [162](#)
- subsystems of AutoSupport [161](#)
- support for AutoSupport, mail host [164](#)
- Syslog Translator
 - getting error message information by using [287](#)
- system
 - date and time, setting [136](#)
 - panics [286](#)
 - password, changing [123](#)
 - rebooting, from the console [90](#)
 - Remote Shell access (RSH) [56](#)
 - time zone, setting [138](#)
- system access
 - starting a Telnet session [53](#)
- system console sessions
 - RLM CLI and [227](#)
 - SP CLI and [197](#)
- system files, Data ONTAP
 - .bplustvoc_internal [132](#)
 - .vtoc_internal [132](#)
- system health
 - See health monitoring
- system information
 - using perfstat to get [251](#)
- System Manager
 - about [61](#)
 - supported Data ONTAP versions [61](#)
 - tasks you can perform from [61](#)
- system operation priorities
 - assigning using FlexShare [257](#)
- system panics
 - considerations for managing core dump files [139](#)
- system prompts
 - rebooting the storage system at the [87](#)
- system reboots
 - automatic technical support notification [141](#)
- system time
 - Data ONTAP options for managing [137](#)
 - synchronizing [137](#)
- systems
 - accessing from Windows client with RSH [58](#)
 - accessing with the remote management device [34](#)
 - allowing only secure access to storage [69](#)
 - components of storage [12](#)
 - disabling root account access to storage [98](#)
 - displaying configuration information for storage [236](#)
 - internal components of storage, defined [12](#)

- management, introduction [20](#)
- methods for administering [28](#)
- methods for administering storage [22](#)
- recovering from a corrupted image of the boot device [88](#)
- remotely rebooting the storage [88, 90](#)
- slots and serial ports of storage [13](#)
- SP admin privilege level commands for managing the [198](#)
- SP advanced privilege level commands for managing the [201](#)
- synchronizing time for [137](#)
- using the RLM to remotely manage [215](#)
- using the SP to remotely manage [189](#)
- ways to boot storage [86](#)

systemshell
uses of [128](#)

T

technical support
automatic notification upon system reboots [141](#)

Telnet
configuring a timeout period [55](#)
controlling the timeout period [55](#)
limiting access to [65](#)
session rules for console, SSH-interactive, and [29](#)
starting a session [53](#)
terminating a session [54](#)

time
Data ONTAP options for managing system [137](#)
setting storage system date and [136](#)

TLS
disabling or enabling [51](#)

Transport Layer Security
See TLS

Transport Layer Security (TLS) protocol [46](#)

trigger events
AutoSupport subsystems [161](#)

troubleshooting
considerations for managing core dump files [139](#)
delivery status of AutoSupport messages [177](#)
mail host [180](#)
storage systems with RLM commands [231](#)
storage systems with SP commands [206](#)
using systemshell and diagnostic account for [128](#)

U

updates

- methods of managing SP firmware [207](#)
- user account, changing password for [124](#)
- user capabilities
how assigned [96](#)
- useradmin
examples [118](#)
naming requirements [96](#)
- users
changing passwords [124](#)
creating and assigning to groups [99](#)
creation examples [118](#)
definition of [95](#)
deleting [117](#)
example, creating with custom capabilities [118](#)
examples of creating [118](#)
listing [112](#)
managing root access [97](#)
modifying capabilities of [102](#)
naming requirement [96](#)

V

V-Series systems
root volumes, introduction to [71](#)

version
checking Data ONTAP [89](#)

volume information
displaying [239](#)

volume priorities
assigning using FlexShare [256](#)
removing using FlexShare [259](#)

volumes
recommendations for root [72](#)
sizing considerations for root FlexVol [73](#)

W

WAFL (Write Anywhere File Layout) [260](#)

WAFL external cache
about [260](#)
compared with Flash Pools [260](#)
disabling [261](#)
displaying configuration [263](#)
displaying usage and access information [264](#)
enabling [261](#)
low-priority user data blocks [262](#)
normal user data blocks [262](#)
rewarming [264](#)
system metadata cache [263](#)

WAFL files

- writing [153](#)
- warnings
 - obsolete domain names [49](#)
- weekly AutoSupport messages [157](#), [161](#)
- Windows
 - domain users, granting access to [101](#)
 - special groups [97](#)
- windows clients
 - accessing storage system from, with RSH [58](#)
- workstations
 - reasons to designate as administration hosts [62](#)
- Write Anywhere File Layout (WAFL) [260](#)
- writing files [153](#)