Data ONTAP™ 7.0.1
Storage Management Guide

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

Part number 210-01111-A0
April 2005

Copyright and trademark information

Copyright information

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Portions Copyright © 2001, Sitraka Inc.

Portions Copyright © 2001, iAnywhere Solutions

Portions Copyright © 2001, i-net software GmbH

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Trademark information

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

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

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

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

Network Appliance NetCache is certified RealSystem compatible.

Table of Contents

Preface xi

Chapter 1 **Introduction to NetApp Storage Architecture**. 1

Understanding NetApp storage architecture 2

Understanding the file system and its storage containers 10

Using volumes from earlier versions of Data ONTAP software 19

Chapter 2 **Quick setup for aggregates and volumes**. 23

Planning your aggregate, volume, and qtree setup 24

Configuring data storage 28

Converting from one type of volume to another 33

Overview of aggregate and volume operations. 34

Chapter 3 **Disk and Storage Subsystem Management** 43

Understanding disks 44

Disk access methods 48

Multipath I/O for Fibre Channel disks 49

Clusters 54

Combined filer head disk shelf access 55

Making disks available 58

Available space on new disks. 60

Adding disks to the system 62

Removing disks. 66

Software-based disk ownership. 71

Re-using disks configured for software-based disk ownership. 76

Disk sanitization 80

Disk and storage subsystem management 92

Changing the state of an adapter 94

Viewing storage subsystem information 96

Chapter 4 **RAID Protection of Data** 105
Understanding RAID groups 106
Predictive disk failure and Rapid RAID Recovery 114
Disk failure and RAID reconstruction with a hot spare disk 115
Disk failure without a hot spare disk 116
Replacing disks in a RAID group 118
Setting RAID type and group size 119
Changing the RAID type for an aggregate 122
Changing the size of RAID groups 127
Controlling the speed of RAID operations 130
Controlling the speed of RAID data reconstruction 131
Controlling the speed of disk scrubbing 132
Controlling the speed of plex resynchronization 133
Controlling the speed of mirror verification 134
Automatic and manual disk scrubs 135
Scheduling an automatic disk scrub 136
Manually running a disk scrub 139
Minimizing media error disruption of RAID reconstructions 142
Handling of media errors during RAID reconstruction 143
Continuous media scrub 144
Disk media error failure thresholds 149
Viewing RAID status 150

Chapter 5 **Aggregate Management** 153
Understanding aggregates 154
Creating aggregates 157
Changing the state of an aggregate 162
Adding disks to aggregates 167
Destroying aggregates 172
Physically moving aggregates 174

Chapter 6 **Volume Management** 177
Traditional and FlexVol volumes 178

Traditional volume operations 181
Creating traditional volumes 182
Physically transporting traditional volumes 187

FlexVol operations 190
Creating FlexVol volumes 191
Resizing FlexVol volumes 195
Cloning FlexVol volumes 197
Displaying a FlexVol's containing aggregate 203

General volume operations 204
Migrating between traditional volumes and FlexVol volumes 205
Managing duplicate volume names 212
Managing volume languages 213
Determining volume status and state 216
Renaming volumes 223
Destroying volumes 224
Increasing the maximum number of files in a volume 226
Reallocating file and volume layout 228

Managing FlexCache volumes 229
How FlexCache volumes work 230
Sample FlexCache deployments 236
Creating FlexCache volumes 238
Sizing FlexCache volumes 240
Administering FlexCache volumes 242

Space management for volumes and files 244
Space guarantees 247
Space reservations 252
Fractional reserve 254

Chapter 7 **Qtree Management** 255

Understanding qtrees 256

Understanding qtree creation 258

Creating qtrees 260

Understanding security styles 261

Changing security styles 263

Changing the CIFS oplocks setting 265

Displaying qtree status 268

Displaying qtree access statistics 269

Converting a directory to a qtree 270

Renaming or deleting qtrees 273

Chapter 8 **Quota Management** 275

Understanding quotas 276

When quotas take effect 279

Understanding default quotas 280

Understanding derived quotas 281

How Data ONTAP identifies users for quotas 284

Notification when quotas are exceeded 287

Understanding the /etc/quotas file 288
- Overview of the /etc/quotas file 289
- Fields of the /etc/quotas file 292
- Sample quota entries 298
- Special entries for mapping users 301
- How disk space owned by default users is counted 305

Activating or reinitializing quotas 306

Modifying quotas 309

Deleting quotas 312

Turning quota message logging on or off 314

Effects of qtree changes on quotas 316

Understanding quota reports 318
- Types of quota reports 319
- Overview of the quota report format 320
- Quota report formats 322
- Displaying a quota report 326

Chapter 9 **SnapLock Management** 327

About SnapLock 328

Creating SnapLock volumes 330

Managing the compliance clock 332

Setting volume retention periods 334

Destroying SnapLock volumes and aggregates 337

Managing WORM data 339

Glossary 341

Index 347

Preface

Introduction

This guide describes how to configure, operate, and manage the storage resources of Network Appliance™ storage systems that run Data ONTAP™ 7.0.1 software. It covers all models. This guide focuses on the storage resources, such as disks, RAID groups, plexes, and aggregates, and how file systems, or volumes, are used to organize and manage data.

Audience

This guide is for system administrators who are familiar with operating systems, such as the UNIX®, Windows NT®, Windows 2000®, Windows Server 2003 Software®, or Windows XP® operating systems, that run on the appliance's clients. It also assumes that you are familiar with how to configure the appliance and how Network File System (NFS), Common Internet File System (CIFS), and Hypertext Transport Protocol (HTTP) are used for file sharing or transfers. This guide doesn't cover basic system or network administration topics, such as IP addressing, routing, and network topology.

Terminology

NetApp® storage products (filers, FAS appliances, and NearStore® systems) are all storage systems—also sometimes called filers or storage appliances.

The terms "flexible volumes" and "FlexVol™ volumes" are used interchangeably in Data ONTAP documentation.

This guide uses the term *type* to mean pressing one or more keys on the keyboard. It uses the term *enter* to mean pressing one or more keys and then pressing the Enter key.

Command conventions

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

In examples that illustrate commands executed on a UNIX workstation, this guide uses the command syntax of SunOS 4.1.x. The command syntax and output might differ, depending on your version of UNIX.

Keyboard conventions

When describing key combinations, this guide uses the hyphen (-) to separate individual keys. For example, *Ctrl-D* means pressing the *Control* and *D* keys simultaneously. Also, this guide uses the term *enter* to refer to the key that generates a carriage return, although the key is named “Return” on some keyboards.

Typographic conventions

The following table describes typographic conventions used in this guide.

Convention	Type of information
Italic font	Words or characters that require special attention. Placeholders for information you must supply. For example, if the guide says to enter the `arp -d` *hostname* command, you enter the characters **`arp -d`** followed by the actual name of the host. Book titles in cross-references.
`Monospaced` font	Command and daemon names. Information displayed on the system console or other computer monitors. The contents of files.
`Bold monospaced` font	Words or characters you type. What you type is always shown in lowercase letters, unless you must type it in uppercase letters.

Special messages

This guide contains special messages that are described as follows:

Note

A note contains important information that helps you install or operate the filer efficiently.

Caution

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

Introduction to NetApp Storage Architecture 1

About this chapter

This chapter provides an overview of how you use Data ONTAP 7.0.1 software to organize and manage the data storage resources (disks) that are part of a NetApp® system and the data that resides on those disks.

Topics in this chapter

This chapter discusses the following topics:

- "Understanding NetApp storage architecture" on page 2
- "Understanding the file system and its storage containers" on page 10
- "Using volumes from earlier versions of Data ONTAP software" on page 19

Understanding NetApp storage architecture

About NetApp storage architecture

NetApp architecture refers to how Data ONTAP utilizes NetApp appliances to make data storage resources available to host or client systems and applications. Data ONTAP 7.0 and later versions distinguish between the physical layer of data storage resources and the logical layer that includes the file systems and the data that reside on the physical resources.

The physical layer includes disks, Redundant Array of Independent Disks (RAID) groups they are assigned to, plexes, and aggregates. The logical layer includes volumes, qtrees, Logical Unit Numbers (LUNs), and the files and directories that are stored in them. Data ONTAP also provides Snapshot™ technology to take point-in-time images of volumes and aggregates.

How NetApp systems use disks

NetApp systems use disks from a variety of manufacturers. All new systems use block checksum disks (BCDs) for RAID parity checksums. These disks provide better performance for random reads than zoned checksum disks (ZCDs), which were used in older systems. For more information about disks, see "Understanding disks" on page 44.

How Data ONTAP uses RAID

Data ONTAP organizes disks into RAID groups, which are collections of data and parity disks to provide parity protection. Data ONTAP supports the following RAID types for NetApp appliances (including the R000 series, the F740, the F800 series, the FAS200 series, and the FAS900 series appliances).

- RAID-4: Before Data ONTAP 6.5, RAID-4 was the only RAID protection scheme available for Data ONTAP aggregates. Within its RAID groups, it allots a single disk for holding parity data, which ensures against data loss due to a single disk failure within a group.
- RAID-DP™ technology (DP for double-parity): RAID-DP provides a higher level of RAID protection for Data ONTAP aggregates. Within its RAID groups, it allots one disk for holding parity data and one disk for holding double-parity data. Double-parity protection ensures against data loss due to a double disk failure within a group.

NetApp gFiler™ gateways support storage systems that use RAID-1, RAID-5, and RAID-10 levels, although the gFilers do not themselves use RAID-1, RAID5, or RAID-10. For information about gFilers and how they support RAID types, see the *gFiler Gateway Series Planning Guide*.

Choosing the right size and the protection level for a RAID group depends on the kind of data you intend to store on the disks in that RAID group. For more information about RAID groups, see “Understanding RAID groups” on page 106.

What a plex is

A *plex* is a collection of one or more RAID groups that together provide the storage for one or more Write Anywhere File Layout (WAFL®) file system *volumes*. Data ONTAP uses plexes as the unit of RAID-level mirroring when SyncMirror is enabled. All RAID groups in one plex are of the same type, but may have a different number of disks.

What an aggregate is

An *aggregate* is a collection of one or two plexes, depending on whether you want to take advantage of RAID-level mirroring. If the aggregate is unmirrored, it contains a single plex. If the SyncMirror® feature is licensed and enabled, you can add a second plex to any aggregate, which serves as a RAID-level mirror for the first plex in the aggregate.

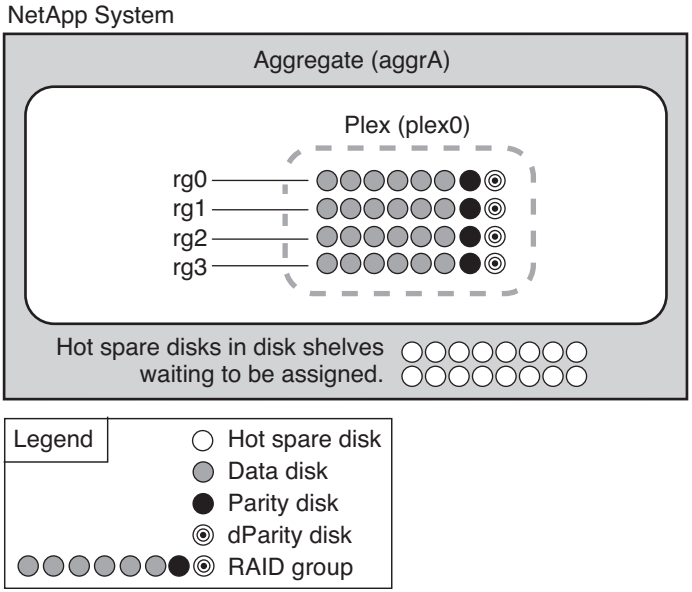
When you create an aggregate, Data ONTAP assigns data disks and parity disks to RAID groups, depending on the options you choose, such as the size of the RAID group (based on the number of disks to be assigned to it) or the level of RAID protection.

You use aggregates to manage plexes and RAID groups because these entities only exist as part of an aggregate. You can increase the usable space in an aggregate by adding disks to existing RAID groups or by adding new RAID groups. Once you’ve added disks to an aggregate, you cannot remove them to reduce storage space without first destroying the aggregate.

If the SyncMirror feature is licensed and enabled, you can convert an unmirrored aggregate to a mirrored aggregate and vice versa without any downtime.

An unmirrored aggregate: Consists of one plex, automatically named by Data ONTAP as *plex0*. This is the default configuration. In the following diagram, the unmirrored aggregate, arbitrarily named *aggrA* by the user, consists of one plex, which is made up of three double-parity RAID groups, automatically named *rg0*, *rg1*, and *rg2* by Data ONTAP.

Notice that RAID-DP requires that both a parity disk and a double parity disk be in each RAID group. In addition to the disks that have been assigned to RAID groups, there are sixteen hot spare disks in one pool of disks waiting to be assigned.

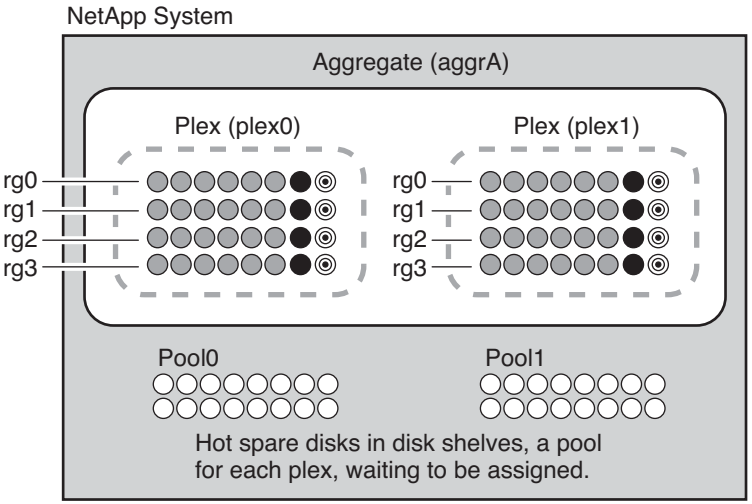


A mirrored aggregate: Consists of two plexes, which provides an even higher level of data redundancy via RAID-level mirroring. For an aggregate to be enabled for mirroring, the appliance must have a SyncMirror license for syncmirror_local or cluster_remote installed and enabled, and the appliance's disk configuration must support RAID-level mirroring.

When you enable SyncMirror, Data ONTAP divides all the disks into two disk pools to ensure a single failures does not affect disks in both pools. This allows the creation of mirrored aggregates. Mirrored aggregates have two plexes, plex0 and plex1. Data ONTAP uses disks from one pool to create plex0 and another pool to create plex1. This provides fault isolation of plexes. A failure that affects one plex will not affect the other plex.

The plexes are physically separated (each plex has its own RAID groups and its own disk pool), and the plexes are updated simultaneously during normal operation. This provides added protection against data loss if there is a double-disk failure or a loss of disk connectivity, because the unaffected plex continues to serve data while you fix the cause of the failure. Once the plex that had a problem is fixed, you can resynchronize the two plexes and reestablish the mirror relationship.

In the following diagram, SyncMirror is enabled, so plex0 has been copied and automatically named *plex1* by Data ONTAP. Notice that plex0 and plex1 contain copies of one or more file systems and that the hot spare disks have been separated into two pools, Pool0 and Pool1.

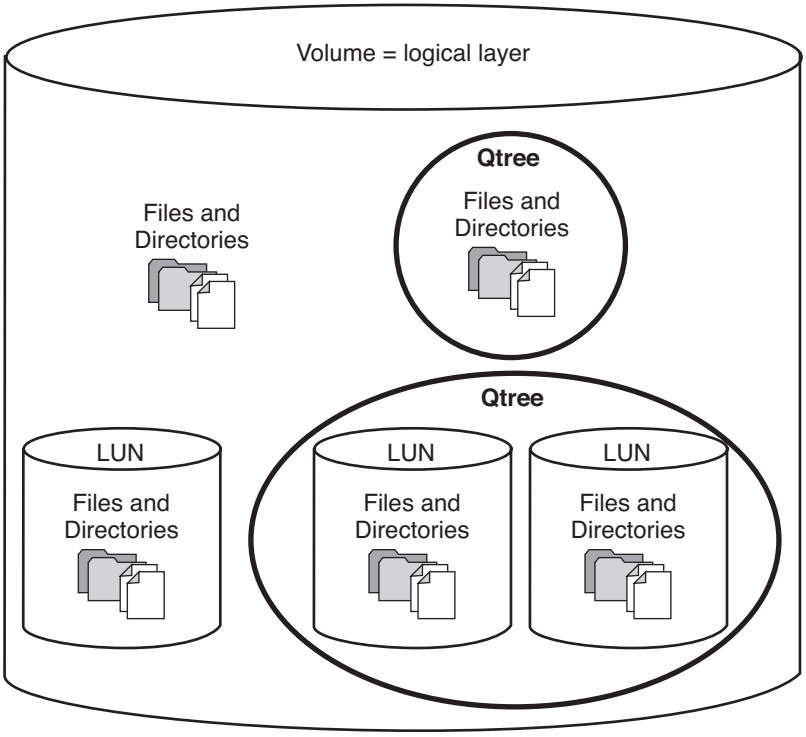


For more information about aggregates, see "Understanding aggregates" on page 154.

What volumes are

Volumes hold user data that is accessible via one or more of the access protocols supported by Data ONTAP, including Network File System (NFS), Common Internet File System (CIFS), HyperText Transfer Protocol (HTTP), Web-based Distributed Authoring and Versioning (WebDAV), Direct Access File System (DAFS), Fibre Channel Protocol (FCP), and Internet SCSI (iSCSI). A volume can include files (which are the smallest units of data storage that hold user- and system-generated data) and, optionally, directories and qtrees in a Network Attached Storage (NAS) environment, and also LUNs in a Storage Area Network (SAN) environment.

The following diagram shows how you can use volumes, qtrees, and LUNs to store files and directories. A volume is the most inclusive of the logical containers. It can store files and directories, qtrees, and LUNs. You can use qtrees to organize files and directories, as well as LUNs. You can use LUNs to serve as virtual disks in SAN environments to store files and directories.



How aggregates provide storage for volumes

Each volume depends on its containing aggregate for all its physical storage. The way a volume is associated with its containing aggregate depends on whether the volume is a *traditional* volume or a *flexible* volume (also referred to as a FlexVol).

Traditional volume: A traditional volume is contained by a single, dedicated, aggregate. A traditional volume is tightly coupled with its containing aggregate. The only way to increase the size of a traditional volume is to add entire disks to its containing aggregate. It is impossible to decrease the size of a traditional volume.

The smallest possible traditional volume must occupy all of two disks (for RAID-4) or three disks (for RAID-DP). Thus, the minimum size of a traditional volume depends on the size and number of disks used to create the traditional volume.

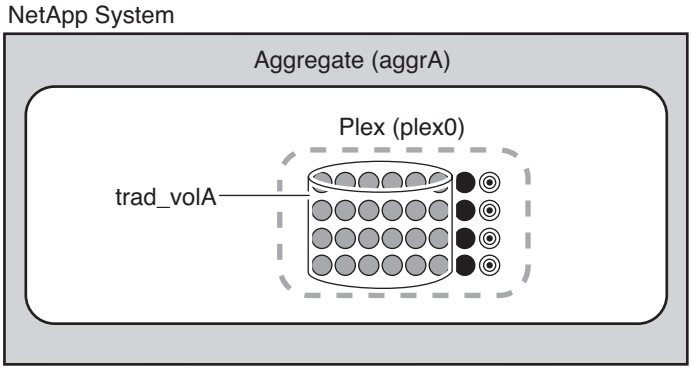
No other volume can use the storage associated with a traditional volume's containing aggregate.

When you create a traditional volume, Data ONTAP creates its underlying containing aggregate based on the parameters you choose with the `vol create` command or with the FilerView® Volume Wizard. Once created, you can manage the traditional volume's containing aggregate with the `aggr` command. You can also use FilerView to perform some management tasks.

The aggregate portion of each traditional volume is assigned its own pool of disks that are used to create its RAID groups, which are then organized into one or two plexes. Because traditional volumes are defined by their own set of disks and RAID groups, they exist outside of and independently of any other aggregates that might be defined on the NetApp appliance.

The following diagram illustrates how a traditional volume, trad_volA, is tightly coupled to its containing aggregate. When volA was created, its size was determined by the amount of disk space requested, the number of disks and their capacity to be used, or a list of disks to be used.

A traditional volume with its tightly-coupled containing aggregate



Flexible volume (FlexVol): A FlexVol is loosely coupled with its containing aggregate. Because the volume is managed separately from the aggregate, FlexVol volumes give you a lot more options for managing the size of the volume. FlexVol volumes provide the following advantages:

- You can create FlexVol volumes in an aggregate nearly instantaneously. They can be as small as 20 MB and as large as the volume capacity that is supported for your platform. For information on the maximum raw volume size supported on NetApp appliances, see the *System Configuration Guide* on the NetApp on the Web™ site (NOW™) at http://netapp.now.com/.

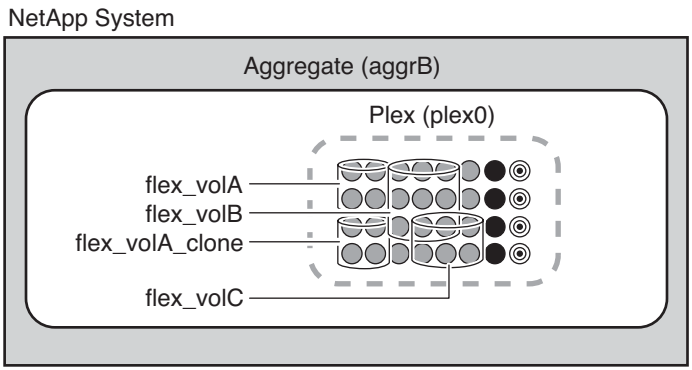
  These volumes stripe their data across all the disks and RAID groups in their containing aggregate.

- You can increase and *decrease* the size of a FlexVol in small increments (as small as 4 KB), nearly instantaneously.
- You can increase the size of a FlexVol to be *larger* than its containing aggregate, which is referred to as aggregate overcommitment. For information about this feature, see "Aggregate overcommitment" on page 250.
- You can clone a FlexVol, which is then referred to as a FlexClone™ volume. For information about this feature, see "Cloning FlexVol volumes" on page 197.

A FlexVol can share its containing aggregate with other FlexVol volumes. Thus, a single aggregate is the shared source of all the storage used by the FlexVol volumes it contains.

In the following diagram, aggrB contains four FlexVol volumes of varying sizes. Note that one of the flexible volumes is a FlexClone.

Flexible volumes with their loosely-coupled containing aggregate



Traditional volumes and FlexVol volumes can co-exist

You can create traditional volumes and FlexVol volumes on the same appliance, up to the maximum number of volumes allowed. For information about maximum limits, see "Maximum numbers of volumes" on page 25.

What snapshots are

A snapshot is a space-efficient, point-in-time image of the data in a volume or an aggregate. Snapshots are used for such purposes as backup and error recovery.

Data ONTAP automatically creates and deletes snapshots of data in volumes to support commands related to Snapshot technology. Data ONTAP also automatically creates snapshots of aggregates to support commands related to the SnapMirror® software, which provides volume-level mirroring. For example, Data ONTAP uses snapshots when data in two plexes of a mirrored aggregate need to be resynchronized.

You can accept the automatic snapshot schedule, or modify it. You can also create one or more snapshots at any time. For more information about snapshots, plexes, and SyncMirror, see the *Data Protection Online Backup and Recovery Guide*.

Understanding the file system and its storage containers

How volumes are used

A volume is a logical file system whose structure is made visible to users when you export the volume to a UNIX host through an NFS mount or to a Windows host through a CIFS share. You assign the following attributes to every volume, traditional or flexible:

- The size of the volume
- A security style, which determines whether a volume can contain files that use UNIX security, files that use NT file system (NTFS) file security, or both types of files
- Whether the volume uses CIFS oplocks (opportunistic locks)
- The type of language supported
- The level of space guarantees
- Disk space and file limits (quotas)
- A snapshot schedule
- Whether the volume is designated as a SnapLock™ volume
- Whether the volume is a *root* volume

  With all new appliances, Data ONTAP is installed at the factory with a root volume already configured. The root volume is named vol0 by default.

  Note

  Even though you can rename any volume at any time, NetApp recommends that you do not rename the root volume.

  - If the root volume is a FlexVol, its containing aggregate is named aggr0 by default.
  - If the root volume is a traditional volume, its containing aggregate is also named vol0 by default. In Data ONTAP 7.0 and later versions, a traditional volume and its containing aggregate always have the same name.

  The root volume contains the appliance's configuration files, the /etc/rc file, which includes startup commands, and log files. You use the root volume to set up and maintain the configuration files. Only one root volume is allowed on an appliance. Since the root volume contains log files, NetApp recommends that your root volume spans four to six disks to handle the increased traffic.

For more information about volumes, see Chapter 6, "Volume Management," on page 177.

How qtrees are used

A qtree is a logically-defined file system that exists as a special top-level subdirectory of the root directory within a volume. You can specify the following features for a qtree.

- A security style like that of volumes
- Whether the qtree uses CIFS oplocks
- Whether the qtree has quotas (disk space and file limits)

  Using quotas enables you to manage storage resources on a per user, user group, or per project status. In this way, you can customize areas for projects and keep users and projects from monopolizing resources.

For more information about qtrees, see Chapter 7, "Qtree Management," on page 255.

How LUNs are used

NetApp storage architecture utilizes two types of LUNs:

- In SAN environments, NetApp systems are targets that have storage target devices, which are referred to as LUNs. With Data ONTAP, you configure NetApp appliances by creating traditional volumes to store LUNs or by creating aggregates to contain FlexVol volumes to store LUNs.

  LUNs created on any NetApp system (both appliances and gFiler gateways) in a SAN environment are used as targets for external storage that is accessible from initiators, or hosts. You use these LUNs to store files and directories accessible through a UNIX or Windows host via FCP or iSCSI.

  For more information about LUNs and how to use them, see the *Block Access Management Guide for FCP* or the *Block Access Management Guide for iSCSI*.
- With gFiler gateways, LUNs are also used for external storage. They are created on the storage subsystems and are available for a gFiler or non-gFiler host to read data from or write data to.

  With gFilers, LUNs provisioned from their external RAID arrays play the role of disks on a NetApp system so that the LUNs themselves provide their own internal RAID array. For more information, see the *gFiler Gateway Series Planning Guide*.

How files are used

Files are the smallest unit of data management. Data ONTAP and application software create system-generated files, and you or your users create data files. You and your users can also create directories in which to store files. You create volumes in which to store files and directories. You create qtrees to organize your volumes. You manage file properties by managing the volume or qtree in which the file or its directory is stored.

How to use NetApp storage resources

The following table describes the storage resources available with NetApp Data ONTAP 7.0 and later versions and how you use them.

Storage Container	Description	How to Use
Disk	Advanced Technology Attachment (ATA) or Fibre Channel disks are used, depending on the platform. Some disk management functions are appliance-specific. For example, for FAS270 appliances, at the boot menu, you specify how many disks the system will own. For R100, R150, R200, FAS800 series, and FAS900 series appliances, Data ONTAP automatically assigns ownership of disks at startup. For gFiler platforms, you must assign ownership of disks.	Once disks are assigned to an appliance, you can choose one of the following methods to assign disks to each RAID group when you create an aggregate: ◆ You provide a list of disks. ◆ You specify a number of disks and let Data ONTAP assign the disks automatically. ◆ You specify the number of disks together with the disk size and/or speed, and let Data ONTAP assign the disks automatically. The following disk-level operations are described in Chapter 3, “Disk and Storage Subsystem Management,” on page 43. ◆ Adding disks ◆ Assigning disk ownership ◆ Configuring Multipath I/O ◆ Sanitizing disks

Storage Container	Description	How to Use
RAID group	Data ONTAP supports RAID-4 and RAID-DP for all filer platforms, and RAID-0 for gFiler platforms. The number of disks that each RAID level uses by default is platform specific.	The smallest RAID group for RAID-4 is two disks (one data and one parity disk); for RAID-DP, it's three (one data and two parity disks). For information about performance, see "Larger versus smaller RAID groups" on page 112. You manage RAID groups with the `aggr` command and FilerView. (For backward compatibility, you can also use the `vol` command for traditional volumes.) The following RAID-level operations are described in Chapter 4, "RAID Protection of Data," on page 105. ◆ Conducting pre-fail disk copy ◆ Conducting preventive disk-media-level scrubs ◆ Configuring RAID group size ◆ Conducting RAID level disk scrubs ◆ Configuring RAID type
Plex	Data ONTAP uses plexes to organize file systems for RAID-level mirroring.	You can ◆ Configure and manage SyncMirror backup replication. For more information, see the *Data Protection Online Backup and Recovery Guide.* ◆ Split an aggregate in a SyncMirror relationship into its component plexes. ◆ View the status of plexes.

Storage Container	Description	How to Use
Aggregate	Consists of one or two plexes. A loosely coupled container for one or more FlexVol volumes. A tightly coupled container for exactly one traditional volume.	You use aggregates to manage disks, RAID groups, and plexes. You can create aggregates implicitly by using the `vol` command to create traditional volumes, explicitly by using the new `aggr` command, or by using the FilerView browser interface The following aggregate-level operations are described in Chapter 5, "Aggregate Management," on page 153. ◆ Adding disks to an aggregate ◆ Changing the state of an aggregate ◆ Creating an aggregate ◆ Destroying an aggregate ◆ Mirroring an aggregate ◆ Physically transporting an aggregate ◆ Renaming an aggregate

Storage Container	Description	How to Use
Volume (common attributes)	Both traditional and FlexVol volumes contain user-visible directories and files, and they can also contain qtrees and LUNs.	You can apply the following volume operations to both FlexVol and traditional volumes. The operations are also described in "General volume operations" on page 204. ◆ Changing the language option for a volume ◆ Changing the state of a volume ◆ Changing the root volume ◆ Destroying volumes ◆ Exporting a volume using CIFS, NFS, and other protocols ◆ Increasing the maximum number of files in a volume ◆ Renaming volumes The following operations are described in the *Data Protection Online Backup and Recovery Guide*. ◆ Implementing SnapMirror ◆ Taking snapshots of volumes The following operation is described later in this guide. ◆ Implementing the SnapLock™ feature

Storage Container	Description	How to Use
Flexible volume (FlexVol)	A logical file system of user data, metadata, and snapshots that is loosely coupled to its containing aggregate. All FlexVol volumes share the underlying aggregate's disk array, RAID group, and plex configurations. Multiple FlexVol volumes can be contained within the same aggregate, sharing its disks, RAID groups, and plexes. FlexVol volumes can be modified and sized independently of their containing aggregate.	You can create FlexVol volumes after you have created the aggregates to contain them. You can increase and decrease the size of a FlexVol by adding or removing space in increments of 4 KB, and you can clone FlexVol volumes. The following FlexVol-level operations are described in "FlexVol operations" on page 190. ◆ Creating FlexVol volumes ◆ Resizing FlexVol volumes ◆ Cloning FlexVol volumes ◆ Displaying the FlexVol's containing aggregate
Traditional volume	A logical file system of user data, metadata and snapshots that is tightly coupled to its containing aggregate. Exactly one traditional volume can exist within its containing aggregate, with the two entities becoming indistinguishable and functioning as a single unit. Traditional volumes are identical to volumes created with earlier than 7.0 versions of Data ONTAP. If you upgrade to Data ONTAP 7.0 and later versions, existing volumes are preserved as traditional volumes.	You can create traditional volumes, physically transport them, and increase them by adding disks. For information about creating and transporting traditional volumes, see "Traditional volume operations" on page 181. For information about increasing the size of a traditional volume, see "Adding disks to aggregates" on page 167.

Storage Container	Description	How to Use
Qtree	An optional, logically defined file system that you can create at any time within a volume. It is a subdirectory of the root directory of a volume. You store directories, files, and LUNs in qtrees. You can create up to 4,995 qtrees per volume.	You use qtrees as logical subdirectories to perform file system configuration and maintenance operations. Within a qtree, you can assign limits to the space that can be consumed and the number of files that can be present (through quotas) to users on a per-qtree basis, define security styles, and enable CIFS opportunity locks (oplocks). The following qtree-level operations are described in Chapter 7, "Qtree Management," on page 255. ◆ Creating qtrees ◆ Changing the security style of a volume or a qtree ◆ Changing the CIFS oplocks setting ◆ Displaying qtree status ◆ Making a directory into a qtree or a qtree into a directory ◆ Deleting qtrees The following qtree-level operations related to configuring usage quotas are described in Chapter 8, "Quota Management," on page 275. ◆ Editing the /etc/quotas file ◆ Turning quotas on or off ◆ Turning quota-message logging on or off ◆ Creating or resizing quotas ◆ Deleting quotas ◆ Reading quota reports

Storage Container	Description	How to Use
LUN (in a SAN environment)	Logical Unit Number; it is a logical unit of storage, which is identified by a number by the initiator accessing its data in a SAN environment. A LUN is a file that appears as a disk drive to the initiator.	You create LUNs within volumes and specify their sizes. For more information about LUNs, see the *Block Access Management Guide for FCP* or *Block Access Management Guide for iSCSI*.
LUN (with gFiler gateways)	An area on the storage subsystem that is available for a gFiler or non-gFiler host to read data from or write data to. The gFiler can virtualize the storage attached to it and serve the storage up as LUNs to customers outside the gFiler (for example, through iSCSI). These LUNs are referred to as gFiler-served LUNs. The clients are unaware of where such a LUN is stored.	See the *gFiler Gateway Series Planning Guide* and the gFiler *Integration Guide* for your storage subsystem for specific information about LUNs and how to use them for your platform
File	Files contain system-generated or user-created data. Files are the smallest unit of data management. Users organize files into directories. As a system administrator, you organize directories into volumes.	Configuring file space reservation is described in Chapter 6, "Volume Management," on page 177.

Using volumes from earlier versions of Data ONTAP software

Upgrading to Data ONTAP 7.0 or later

If you are upgrading to Data ONTAP 7.0 or later software from an earlier version, your existing volumes are preserved as traditional volumes. Your volumes and data remain unchanged, and the commands you used to manage your volumes and data are still supported for backward compatibility.

As you learn more about FlexVol volumes, you might want to migrate your data from traditional volumes to FlexVol volumes. For information about migrating traditional volumes to FlexVol volumes, see "Migrating between traditional volumes and FlexVol volumes" on page 205.

Using traditional volumes

With traditional volumes, you can use the new `aggr` and `aggr options` commands or FilerView to manage its containing aggregate. For backward compatibility, you can also use the `vol` and the `vol options` commands to manage the traditional volume's containing aggregate.

The following table describes how to create and manage traditional volumes using either the `aggr` or the `vol` commands, and FilerView, depending on whether you are managing the physical or logical layers of that volume.

Traditional volume task, using FilerView, if available	**Using the `aggr` command**	**Using the `vol` command**

Create a volume
In FilerView:
Volumes > Add
	Not applicable.	`aggr create` *`trad_vol`* `-v -m {`*`disk-list`* `|` *`size`*`}`
		Creates a traditional volume and defines a set of disks to include in that volume or defines the size of the volume.
		The -v option designates that `trad_vol` is a traditional volume.
		Use `-m` to enable SyncMirror.
		For backward compatibility, use:
		`vol create` *`trad_vol`* `-m { `*`disk-list`* `|` *`size`* `}`

Traditional volume task, using FilerView, if available	Using the `aggr` command	Using the `vol` command
Add disks In FilerView: Volumes > Manage. Click the *trad_vol* name you want to add disks to. The Volume Properties page appears. Click Add disks. The Volume Wizard appears.	`aggr add` *`trad_vol disks`*	For backward compatibility: `vol add` *`trad_vol disks`*

Create a SyncMirror replica
 	
In FilerView:
 	
For new aggregates:
Aggregates > Add
 	
For existing aggregates:
Aggregates > Manage.
Click *trad_vol*.
The Aggregate properties page appears.
Click Mirror
Click OK
	`aggr mirror`
	
	`aggr split`
	
	`aggr verify`
		For backward compatibility:
		
		`vol mirror`
		
		`vol split`
		
		`vol verify`

Set the root volume option This option can be used on only one volume per appliance. For more information on root volumes, see “How volumes are used” on page 10.	Not applicable.	`vol options` *`trad_vol`* `root` If the root option is set on a traditional volume, that volume becomes the root volume for the appliance on the next reboot.

Traditional volume task, using FilerView, if available	**Using the `aggr` command**	**Using the `vol` command**
Set RAID options In FilerView: For new aggregates: Aggregates > Add For existing aggregates: Aggregates > Manage. Click *trad_vol*. Click Modify	`aggr options` *`trad_vol`* `{` *`raidsize`* `|` *`raidtype`* `}`	For backward compatibility: `vol options` *`trad_vol`* `{` *`raidsize`* `|` *`raidtype`* `}`
Set up SnapLock volume	Not applicable.	`aggr create` *`trad_vol`* `-r -L` *`disk-list`* Use `vol create trad_vol -r -L` *`disk-list`* for backward compatibility.

RAID level scrub
In FilerView:
Aggregates >
Configure RAID
	`aggr scrub start`
	`aggr scrub suspend`
	`aggr scrub stop`
	`aggr scrub resume`
	`aggr scrub status`
	Manages RAID-level error scrubbing of the disks.
	See “Automatic and manual disk scrubs” on page 135.
		For backward compatibility:
		`vol scrub start`
		`vol scrub suspend`
		`vol scrub stop`
		`vol scrub resume`
		`vol scrub status`

Media level scrub	`aggr media_scrub` *`t_vol`* Manages media error scrubbing of disks in the traditional volume. See “Continuous media scrub” on page 144.	For backward compatibility: `vol media_scrub` *`t_vol`*

Quick setup for aggregates and volumes 2

About this chapter

This chapter provides the information you need to plan and create aggregates and volumes.

After initial setup of your appliance's disk groups and file systems, you can manage or modify them using information in other chapters.

Topics in this chapter

This chapter discusses the following topics:

- "Planning your aggregate, volume, and qtree setup" on page 24
- "Configuring data storage" on page 28
- "Converting from one type of volume to another" on page 33
- "Overview of aggregate and volume operations" on page 34

Planning your aggregate, volume, and qtree setup

Planning considerations

How you plan to create your aggregates and FlexVol volumes, traditional volumes, qtrees, or LUNs depends on your requirements and whether your new version of Data ONTAP is a new installation or an upgrade from Data ONTAP 6.5.*x* or earlier. For information about upgrading a NetApp appliance, see the *Data ONTAP 7.0.1 Upgrade Guide*.

Considerations when planning aggregates

For new appliances: If you purchased a new appliance with Data ONTAP 7.0 or later installed, the root FlexVol (vol0) and its containing aggregate (aggr0) are already configured.

The remaining disks on the appliance are all unallocated. You can create any combination of aggregates with FlexVol volumes, traditional volumes, qtrees, and LUNs, according to your needs.

Maximizing storage: To maximize the storage capacity of your filer per volume, configure large aggregates containing multiple FlexVol volumes. Because multiple FlexVol volumes within the same aggregate share the same RAID parity disk resources, more of your disks are available for data storage.

SyncMirror Replication: You can set up a RAID-level mirrored aggregate to contain volumes whose users require guaranteed SyncMirror data protection and access. SyncMirror replicates the volumes in plex0 to plex1. The disks used to store the second plex can be up to 30 km away if you use MetroCluster. If you set up SyncMirror replication, plan to allocate double the disks that you would otherwise need for the aggregate to support your users.

All volumes contained in a mirrored aggregate are in a SyncMirror relationship, and all new volumes created within the mirrored aggregate inherit this feature. For more information on configuring and managing SyncMirror replication, see the *Data ONTAP Online Backup and Recovery Guide*.

Size of RAID groups: When you create an aggregate, you can control the size of a RAID group. Generally, larger RAID groups maximize your data storage space by providing a greater ratio of data disks to parity disks. For information on RAID group size guidelines, see “Larger versus smaller RAID groups” on page 112.

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

For more information on RAID-4 and RAID-DP, see "Types of RAID protection" on page 106.

Considerations when planning volumes

Root volume sharing: When NetApp technicians install Data ONTAP on your NetApp system, they create a root volume named *vol0*. The root volume is a FlexVol, so you can resize it. For information about resizing FlexVol volumes, see "Resizing FlexVol volumes" on page 195.

SnapLock volume: The SnapLock feature enables you to keep a permanent snapshot by writing new data once to disks and then preventing the removal or modification of that data. You can create and configure a special traditional volume to provide this type of access, or you can create an aggregate to contain FlexVol volumes that provide this type of access. If an aggregate is enabled for SnapLock, all of the FlexVol volumes that it contains have mandatory SnapLock protection. For more information, see "About SnapLock" on page 328.

Data sanitization: Disk sanitization is a Data ONTAP feature that enables you to erase sensitive data from filer disks beyond practical means of physical recovery. Because data sanitization is carried out on the entire set of disks in an aggregate, configuring smaller aggregates to hold sensitive data that requires sanitization minimizes the time and disruption that sanitization operations entail.You can create smaller aggregates and traditional volumes whose data you might have reason to sanitize at periodic intervals. For more information, see "Disk sanitization" on page 80.

Maximum numbers of aggregates: You can create up to 100 aggregates per appliance, regardless of whether the aggregates contain FlexVol or traditional volumes.

You can use the `aggr status` command or FilerView (by viewing the System Status window) to see how many aggregates exist. With this information, you can determine how many more aggregates you can create on the appliance, depending on available capacity. For more information about FilerView, see the *System Administration Guide*.

Maximum numbers of volumes: You can create up to 200 volumes per appliance. However you can only create up to 100 traditional volumes because of the 100 aggregates per appliance limit. You can use the `vol status` command or FilerView (Volumes > Manage > Filter by) to see how many volumes exist, and

whether they are FlexVol or traditional volumes. With this information, you can determine how many more volumes you can create on that appliance, depending on available capacity.

Consider the following example. Assume you create:

- Ten traditional volumes. Each has exactly one containing aggregate.
- Twenty aggregates, and you then create four FlexVol volumes in each aggregate, for a total of eighty FlexVol volumes.

You now have a total of:

- Thirty aggregates (ten from the traditional volumes, plus the twenty created to hold the FlexVol volumes)
- Ninety volumes (ten traditional and eighty FlexVol) on the appliance

Thus, the appliance is well under the maximum limits for either aggregates or volumes.

If you have a combination of FlexVol and traditional volumes, the 100-maximum limit of aggregates still applies. If you need more than 200 user-visible file systems, you can create qtrees within the volumes.

Considerations for FlexVol volumes

When planning the setup of your FlexVol volumes within an aggregate, consider the following issues.

General Deployment: FlexVol volumes have different best practices, optimal configurations, and performance characteristics compared to traditional volumes. Make sure you understand these differences and deploy the configuration that is optimal for your environment.

For information about deploying a storage solution with FlexVol volumes, including migration and performance considerations, see the technical report *Introduction to Data ONTAP Release 7G* (available from the NetApp Library at http://www.netapp.com/tech_library/ftp/3356.pdf).

FlexVol space guarantee: Setting a maximum volume size does not guarantee that the volume will have that space available if the aggregate space is oversubscribed. As you plan the size of your aggregate and the maximum size of your FlexVol volumes, you can choose to overcommit space if you are sure that the actual storage space used by your volumes will never exceed the physical data storage capacity that you have configured for your aggregate. This is called aggregate overcommitment. For more information, see “Aggregate overcommitment” on page 250.

Volume language: During volume creation you can specify the language character set to be used.

Backup: You can size your FlexVol volumes for convenient volume-wide data backup through SnapMirror, SnapVault™, and Volume Copy features. For more information, see the *Data ONTAP Online Backup and Recovery Guide*.

Volume cloning: Many database programs enable data cloning, that is, the efficient copying of data for the purpose of manipulation and projection operations. This is efficient because Data ONTAP allows you to create a duplicate of a volume by having the original volume and clone volume share the same disk space for storing unchanged data. For more information, see "Cloning FlexVol volumes" on page 197.

Considerations for traditional volumes

Upgrading: If you upgrade to Data ONTAP 7.0 or later from a previous version, the upgrade program preserves each of your existing volumes as traditional volumes.

Disk portability: You can create traditional volumes and aggregates whose disks you intend to physically transport from one appliance to another. This ensures that a specified set of physically transported disks will hold all the data associated with a specified volume and only the data associated with that volume. For more information, see "Physically transporting traditional volumes" on page 187.

Considerations when planning qtrees

Within a volume you have the option of creating qtrees to provide another level of logical file systems. Some reasons to consider setting up qtrees include:

Increased granularity: Up to 4,995 qtrees—that is 4,995 virtually independent file systems—are supported per volume. For more information see Chapter 7, "Qtree Management," on page 255.

Sophisticated file and space quotas for users: Qtrees support a sophisticated file and space quota system that you can use to apply soft or hard space usage limits on individual users, or groups of users. For more information see Chapter 8, "Quota Management," on page 275.

Configuring data storage

About configuring data storage

You configure data storage by creating aggregates and Flex volumes, traditional volumes, and LUNs for a SAN environment. You can also use qtrees to partition data in a volume.

You can create up to 100 aggregates per filer. Minimum aggregate size is two disks (one data disk, one parity disk) for RAID-4 or three disks (one data, one parity, and one double parity disk) for RAID-DP. However, NetApp recommends that you configure the size of your RAID groups according to the anticipated load. For more information, see the chapter on system information and performance in the *System Administration Guide*.

Creating aggregates, FlexVol volumes, and qtrees

To create an aggregate and a FlexVol, complete the following steps.

Step	Action
1	(Optional) Determine the free disk resources on your filer by entering the following command: **`aggr status -s`** `-s` displays a listing of the spare disks on the filer. **Result:** Data ONTAP displays a list of the disks that are not allocated to an aggregate. With a new appliance, all disks except those allocated for the root volume's aggregate (explicit for a FlexVol and internal for a traditional volume) will be listed.

Step	Action
2	(Optional) Determine the size of the aggregate, assuming it is aggr0, by entering one of the following commands: For size in kilobytes, enter: **`df -A aggr0`** For size in 4096-byte blocks, enter: **`aggr status -b aggr0`** For size in number of disks, enter: **`aggr status { -d | -r } aggr0`** `-d` displays disk information `-r` displays RAID information **Note** If you want to expand the size of the aggregate, see "Adding disks to an aggregate" on page 168.
3	Create an aggregate by entering the following command: **`aggr create [ -m ] aggr ndisks[@disksize]`** **Example:** `aggr create aggr1 24@72G` **Result:** An aggregate named *aggr1* is created. It consists of 24 72-GB disks. `-m` instructs Data ONTAP to implement SyncMirror. By default, the RAID level is set to RAID-DP. If you use the group size of sixteen, aggr1 consists of two RAID groups, the first group having fourteen data disks, one parity disk, and one double parity disk, and the second group having six data disks, one parity disk, and one double parity disk. If you use the group size of eight, aggr1 consists of three disk RAID groups, each one having six data disks, one parity disk, and one double parity disk.

Step	Action
4	(Optional) Verify the creation of this aggregate by entering the following command: **`aggr status aggr1`**
5	Create a FlexVol in the specified aggregate by entering the following command: **`vol create vol aggr size`** **Example:** `vol create new_vol aggr1 32g` **Result:** The FlexVol new_vol, with a maximum size of 32 GB, is created in the aggregate, aggr1. The default space guarantee setting for FlexVol creation is `volume`. The `vol create` command fails if Data ONTAP cannot guarantee 32 GB of space. To override the default, enter one of the following commands. For information about space guarantees, see "Space guarantees" on page 247. **`vol create vol -s none aggr size`** or **`vol create vol -s file aggr size`**
6	(Optional) To verify the creation of the FlexVol named *new_vol*, enter the following command: **`vol status new_vol -v`**
7	If you want to create additional FlexVol volumes in the same aggregate, use the `vol create` command as described in Step 5. Note the following constraints: ◆ Volumes must be uniquely named across all aggregates within the same filer. If aggregate *aggr1* contains a volume named *volA*, no other aggregate on the filer can contain a volume with the name *volA*. ◆ You can create a maximum of 200 FlexVol volumes in one filer. ◆ Minimum size of a FlexVol is 20 MB.

Step	Action
8	To create qtrees within your volumes, enter the following command: **`qtree create /vol/vol/qtree`** **Example:** `qtree create /vol/new_vol/my_tree` **Result:** The qtree *my_tree* is created within the volume named new_vol. **Note** You can create up to 4,995 qtrees within one volume.
9	(Optional) To verify the creation of the qtree named my_tree, within the volume named new_vol, enter the following command: **`qtree status new_vol -v`**

Why continue using traditional volumes

If you upgrade to Data ONTAP 7.0 or later from a previous version of Data ONTAP, the upgrade program keeps your traditional volumes intact. You might want to maintain your traditional volumes and create additional traditional volumes because some operations are more practical on traditional volumes, such as:

- Performing disk sanitization operations
- Physically transferring volume data from one location to another (which is most easily carried out on small-sized traditional volumes)
- Migrating volumes using the SnapMover® feature

Creating traditional volumes and qtrees

To create a traditional volume, complete the following steps:

Step	Action
1	(Optional) List the aggregates and traditional volumes on your filer by entering the following command: **`aggr status -v`**

Step	Action
2	(Optional) Determine the free disk resources on your filer by entering the following command: **aggr status -s**
3	To create a traditional volume, enter the following command: **aggr create *trad_vol* -v n*disks*[@*disksize*]** **Example:** `aggr create new_tvol -v 16@72g`
4	(Optional) To verify the creation of the traditional volume named new_tvol, enter the following command: **vol status new_tvol -v 16@72g**
5	If you want to create additional traditional volumes, use the `aggr create` command as described in Step 3. Note the following constraints: ◆ All volumes, including traditional volumes, must be uniquely named within the same filer. ◆ You can create a maximum of 100 traditional volumes within one appliance. ◆ Minimum traditional volume size depends on the disk capacity and RAID protection level.
6	To create qtrees within your volume, enter the following command: **qtree create /vol/*vol*/*qtree*** **Example:** `qtree create /vol/new_tvol/users_tree` **Result:** The qtree *users_tree* is created within the new_tvol volume. **Note** You can create up to 4,995 qtrees within one volume.
7	(Optional) To verify the creation of the qtree named users_tree within the new_tvol volume, enter the following command line: **qtree status new_tvol -v**

Converting from one type of volume to another

What converting to another volume type involves

Converting from one type of volume to another is not a single-step procedure. It involves creating a new volume, migrating data from the old volume to the new volume, and verifying that the data migration was successful. You can migrate data from traditional volumes to FlexVol volumes or vice versa. For more information about migrating data, see "Migrating between traditional volumes and FlexVol volumes" on page 205.

When to convert from one type of volume to another

You might want to convert a traditional volume to a FlexVol because

- You upgraded an existing NetApp system that is running an earlier release than Data ONTAP 7.0 or later and you want to convert the traditional root volume to a FlexVol to reduce the amount of disks used to store the system directories and files.
- You purchased a new system but initially created traditional volumes and now you want to
  - Take advantage of FlexVol volumes
  - Take advantage of other advanced features, such as FlexClone volumes
  - Reduce lost capacity due to the number of parity disks associated with traditional volumes
  - Realize performance improvements by being able to increase the number of disks the data in a FlexVol is striped across

You might want to convert a FlexVol to a traditional volume because

- You want to revert to an earlier release of Data ONTAP.

Depending on the number and size of traditional volumes on your NetApp systems, this might require a significant amount of planning, resources, and time.

NetApp offers assistance

NetApp Professional Services staff, including Professional Services Engineers (PSEs) and Professional Services Consultants (PSCs) are trained to assist customers with converting volume types and migrating data, among other services. For more information, contact your local NetApp Sales representative, PSE, or PSC.

Overview of aggregate and volume operations

About aggregate and volume-level operations

The following table provides an overview of the operations you can carry out on an aggregate, a FlexVol, and a traditional volume.

Operation	Aggregate	FlexVol	Traditional volume
Adding disks to an aggregate	`aggr add` *`aggr disks`* Adds disks to the specified aggregate. See "Adding disks to aggregates" on page 167.	Not applicable.	`aggr add` *`trad_vol disks`* Adds disks to the specified traditional volume. See "Adding disks to aggregates" on page 167.
Changing the size of an aggregate	See "Adding disks to the system" on page 62 and "Adding disks to aggregates" on page 167.	Not applicable.	See "Adding disks to the system" on page 62 and "Adding disks to aggregates" on page 167.
Changing the size of a volume	Not applicable	`vol size` *`flex_vol newsize`* Modifies the size of the specified FlexVol. See "Resizing FlexVol volumes" on page 195.	See changing the size of an aggregate.

Operation	Aggregate	FlexVol	Traditional volume

Changing states: online, offline, restricted	`aggr offline aggr`
	`aggr online aggr`
	`aggr restrict aggr`
	Takes the specified aggregate offline, brings it back online, or puts it in a restricted state.
	See "Changing the state of an aggregate" on page 162.
		`vol offline vol`
		`vol online vol`
		`vol restrict vol`
		Takes the specified volume offline, brings it back online (if its containing aggregate is also online), or puts it in a restricted state.
		See "Determining volume status and state" on page 216.
			`aggr offline vol`
			`aggr online vol`
			`aggr restrict vol`
			Takes the specified volume offline, brings it back online, or puts it in a restricted state.
			See "Determining volume status and state" on page 216.

Copying	`aggr copy start src_aggr dest_aggr` Copies the specified aggregate and its FlexVol volumes to a different aggregate on a new set of disks. See the *Data Protection Online Backup and Recovery Guide*.	`vol copy start src_vol dest_vol` Copies the specified source volume and its data content to a destination volume on a new set of disks. The source and destination volumes must be of the same type (either FlexVol or traditional). See the *Data Protection Online Backup and Recovery Guide*.	

Operation	Aggregate	FlexVol	Traditional volume
Creating an aggregate	`aggr create aggr [-f] [-m] [-n] [-t raidtype] [-r raidsize] [-R rpm] [-T disk-type] [-R rpm] [-L] {ndisks[@size] | -d disk1 [disk2 ...] [-d diskn [diskn+1 ... ]]}` Creates a physical aggregate of disks, within which FlexVol volumes can be created. See "Creating aggregates" on page 157.	Not applicable.	See creating a volume.
Creating a volume	Not applicable.	`vol create flex_vol [-l language_code] [ -s none | file | volume ] aggr size` Creates a FlexVol within the specified containing aggregate. See "Creating FlexVol volumes" on page 191.	`aggr create trad_vol -v [-l language_code] [-f] [-n] [-m] [-L] [-t raidtype] [-r raidsize] [-R rpm] {ndisks@size] | -d disk1 [disk2 ...] [-d diskn [diskn+1 ... ]]}` Creates a traditional volume and defines a set of disks to include in that volume. See "Creating traditional volumes" on page 182.

Operation	Aggregate	FlexVol	Traditional volume
Creating a FlexClone	Not applicable.	`vol clone create` *`flex_vol clone_vol`* Creates a clone of the specified FlexVol. See "Cloning FlexVol volumes" on page 197.	Not applicable.
Creating a SnapLock volume	`aggr create` *`aggr`* `-L` *`disk-list`* See "Creating SnapLock aggregates" on page 330.	FlexVol volumes inherit the SnapLock attribute from their containing aggregate. See "Creating SnapLock volumes" on page 330.	`aggr create` *`trad_vol`* `-v -L` *`disk-list`* See "Creating SnapLock traditional volumes" on page 330.

Creating a SyncMirror replica	`aggr mirror`
	`aggr split`
	`aggr verify`
	Creates a SyncMirror replica of the specified aggregate.
	See the *Data Protection Online Backup and Recovery Guide*.
		Not applicable.	`aggr mirror`
			`aggr split`
			`aggr verify`
			Creates a SyncMirror replica of the specified traditional volume.
			See the *Data Protection Online Backup and Recovery Guide*.

Displaying the containing aggregate	Not applicable.	`vol container` *`flex_vol`* Displays the containing aggregate of the specified FlexVol. See "Displaying a FlexVol's containing aggregate" on page 203.	Not applicable.
Displaying the language code	Not applicable	`vol lang [`*`vol`*`]` Displays the FlexVol's language. See "Changing the language for a volume" on page 215.	

Operation	Aggregate	FlexVol	Traditional volume
Displaying a media-level scrub	`aggr media_scrub status [aggr]` Displays media error scrubbing of disks in the aggregate. See "Continuous media scrub" on page 144	Not applicable.	`aggr media_scrub status [aggr]` Displays media error scrubbing of disks in the traditional volume. See "Continuous media scrub" on page 144.
Displaying the status	`aggr status [aggr]` Displays the offline, restricted, or online status of the specified aggregate. Online status is further defined by RAID state, reconstruction, or mirroring conditions. See "Changing the state of an aggregate" on page 162.	`vol status [vol]` Displays the offline, restricted, or online status of the specified volume, and the RAID state of its containing aggregate. See "Determining volume status and state" on page 216.	`aggr status [vol]` Displays the offline, restricted, or online status of the specified volume. Online status is further defined by RAID state, reconstruction, or mirroring conditions. See "Determining volume status and state" on page 216.
Destroying aggregates and volumes	`aggr destroy aggr` Destroys the specified aggregate and returns that aggregate's disks to the filer's pool of hot spare disks. See "Destroying aggregates" on page 172.	`vol destroy flex_vol` Destroys the specified FlexVol and returns space to its containing aggregate. See "Destroying volumes" on page 224.	`aggr destroy trad_vol` Destroys the specified traditional volume and returns that volume's disks to the filer's pool of hot spare disks See "Destroying volumes" on page 224.

Operation	Aggregate	FlexVol	Traditional volume

Performing a RAID-level scrub	`aggr scrub start`
	`aggr scrub suspend`
	`aggr scrub stop`
	`aggr scrub resume`
	`aggr scrub status`
	Manages RAID-level error scrubbing of disks of the aggregate.
	See "Automatic and manual disk scrubs" on page 135.
		Not applicable.	`aggr scrub start`
			`aggr scrub suspend`
			`aggr scrub stop`
			`aggr scrub resume`
			`aggr scrub status`
			Manages RAID-level error scrubbing of disks of the traditional volume.
			See "Automatic and manual disk scrubs" on page 135

Renaming aggregates and volumes	`aggr rename` *`old_name new_name`* Renames the specified aggregate as *new_name*. See "Renaming an aggregate" on page 166.	`vol rename` *`old_name new_name`* Renames the specified flexible volume as *new_name.* See "Renaming volumes" on page 223.	`aggr rename` *`old_name new_name`* Renames the specified traditional volume as *new_name.* See "Renaming volumes" on page 223.
Setting the language code	Not applicable	`vol lang` *`vol language_code`* Sets the FlexVol's language. See "Changing the language for a volume" on page 215.	
Setting the maximum directory size	Not applicable.	`vol option` *`vol`* `maxdirsize` *`size`* *size* specifies the maximum directory size allowed in the specified volume. See "Increasing the maximum number of files in a volume" on page 226.	

Operation	Aggregate	FlexVol	Traditional volume
Setting the RAID options	`aggr options aggr {raidsize|raidtype}` Modifies RAID settings on the specified aggregate. See "Setting RAID type and group size" on page 119 or "Changing the RAID type for an aggregate" on page 122.	Not applicable.	`aggr options trad_vol {raidsize|raidtype}` Modifies RAID settings on the specified traditional volume. See "Setting RAID type and group size" on page 119 or "Changing the RAID type for an aggregate" on page 122.
Setting the root volume	Not applicable.	`vol options flex_vol root`	`vol options trad_vol root`
Setting the UNICODE options	Not applicable.	`vol options vol {convert _ucode | create_ucode} {on|off}` Forces or specifies as default conversion to UNICODE format on the specified volume. For information about UNICODE, see the *System Administration Guide*.	

Configuring volume-level options

The following table provides an overview of the options you can use to configure your aggregates, FlexVol and traditional volumes.

Note

The `option` subcommands you execute remain in effect after the appliance is rebooted, so you do not have to add `aggr options` or `vol options` commands to the /etc/rc file.

Aggregate	FlexVol	Traditional volume
`aggr options aggr [optname optvalue]` Displays the option settings of *aggr*, or sets *optname* to *optvalue*. See the na_aggr man page.	`vol options vol [optname optvalue]` Displays the option settings of *vol*, or sets *optname* to *optvalue*. See the na_vol man page.	
	`convert_ucode on | off`	`convert_ucode on | off`
	`create_ucode on | off`	`create_ucode on | off`
	`fractional_reserve percent`	`fractional_reserve percent`
`fs_size-fixed on | off`	`fs_size-fixed on | off`	`fs_size-fixed on | off`
	`guarantee file | volume | none`	
`ignore_inconsistent on | off`		`ignore_inconsistent on | off`
`lost_write_protect`		
	`maxdirsize number`	`maxdirsize number`
	`minra on | off`	`minra on | off`
	`no_atime_update on | off`	`no_atime_update on | off`
`nosnap on | off`	`nosnap on | off`	`nosnap on | off`
	`nosnapdir on | off`	`nosnapdir on | off`
	`nvfail on | off`	`nvfail on | off`
`raidsize number`		`raidsize number`
`raidtype raid4 | raid_dp | raid0`		`raidtype raid4 | raid_dp | raid0`
`resyncsnaptime number`		`resyncsnaptime number`
`root`	`root`	`root`

Aggregate	FlexVol	Traditional volume
snaplock_compliance (read only)	snaplock_compliance (read only)	snaplock_compliance (read only)
	snaplock_default_period (read only)	snaplock_default_period (read only)
snaplock_enterprise (read only)	snaplock_enterprise (read only)	snaplock_enterprise (read only)
	snaplock_minimum_period	snaplock_minimum_period
	snaplock_maximum_period	snaplock_maximum_period
snapmirrored off	snapmirrored off	snapmirrored off
snapshot_autodelete on | off		
	svo_allow_rman on | off	svo_allow_rman on | off
	svo_checksum on | off	svo_checksum on | off
	svo_enable on | off	svo_enable on | off
	svo_reject_errors	svo_reject_errors

Disk and Storage Subsystem Management 3

About this chapter

This chapter discusses how NetApp systems use disks, how disks are dually connected to systems for redundancy, how to add disks to systems so that Data ONTAP can use them, and how to remove disks from the system. It also discusses how you can check the status on disks and other storage subsystem components connected to your system, including the adapters, hubs, tape devices, and medium changer devices.

For instructions about adding or removing disks from a disk shelf physically or determining the location of a disk in a disk shelf, see your disk shelf documentation or the hardware and service guide for your system.

Topics in this chapter

This chapter discusses the following topics:

- "Understanding disks" on page 44
- "Disk access methods" on page 48
- "Making disks available" on page 58
- "Available space on new disks" on page 60
- "Adding disks to the system" on page 62
- "Removing disks" on page 66
- "Software-based disk ownership" on page 71
- "Re-using disks configured for software-based disk ownership" on page 76
- "Disk sanitization" on page 80
- "Disk and storage subsystem management" on page 92

Understanding disks

Disk checksum and format

All new NetApp systems use block checksum disks (BCDs), which have a disk format of 520 bytes per sector. If you have an older system, it might have zoned checksum disks (ZCDs), which have a disk format of 512 bytes per sector. When you run the `setup` command, Data ONTAP uses the disk checksum type to determine the checksum type of aggregates that you create. For more information about checksum types, see "How Data ONTAP enforces checksum type rules" on page 157.

Supported disk types

Data ONTAP supports the following disks types:

- FC-AL
- SCSI
- ATA (Parallel ATA)

Data ONTAP does not allow mixing disk types in the same aggregate.

When a disk fails, Data ONTAP can only replace it with a disk of the same type. For information about replacing disks, see "Replacing disks in a RAID group" on page 118 or the na_disk man page.

What determines which disks can be used for a NetApp system

Each NetApp system supports specific disk shelves and I/O modules, which determine the types of disks that can be used for the system. For example, the R200 NearStore® appliance supports the DSmk2 AT shelf, which supports the AT-FC I/O module, which supports ATA disks.

The following table describes which systems support which type of disk shelf and I/O module. For more current information, see the *System Configuration Guide* on the NetApp on the Web (NOW) site at http://now.netapp.com/.

Appliance	Disk Shelf Supported	I/O Module Supported
F87	Internal disk shelf	Not applicable.
FAS250	DS14mk2 FC	Not applicable (no expansion shelves).

Appliance	Disk Shelf Supported	I/O Module Supported
F800 series	Fibre Channel StorageShelf FC7, FC8, FC9	Not applicable.
	DS14 DS14mk2 FC	LRC, ESH, ESH2
FAS270	DS14mk2 FC	LRC, ESH2
FAS920, FAS940	Fibre Channel StorageShelf FC7, FC8, FC9	Not applicable.
	DS14, DS14mk2 FC	LRC, ESH, ESH2
FAS960	Fibre Channel StorageShelf FC7, FC8, FC9	Not applicable.
	DS14, DS14mk2 FC	LRC, ESH, ESH2
FAS980	Fibre Channel StorageShelf FC9	Not applicable.
	DS14, DS14mk2 FC	LRC, ESH, ESH2
R100	R1XX disk shelf (with ATA disks)	Not applicable.
R150	R1XX disk shelf (with ATA disks)	Not applicable.
	DS14mk2 AT	AT-FC
R200	DS14mk2 AT	AT-FC

How disks are initially configured at the factory

After NetApp technicians install the Data ONTAP operating system, they configure some or all of the disks, depending on the model:

- For the NearStore models (R100, R150, and R200), the F800 series, and the FAS900 series, they configure all of the disks as spare disks in a pool, and then create a root volume, which contains the system's configuration files and the /etc/rc file.
- For some systems, such as the FAS250 and FAS270 appliances and gFiler gateways, you must assign disk ownership at first boot before you can create aggregates and volumes. For information about setting up your specific system, see the appropriate hardware guide for your system and the *Software Setup Guide*.

The disks are installed with the latest firmware. However, you might need to upgrade disk firmware for FC-AL or SCSI disks when new firmware is offered, or when you upgrade the Data ONTAP software.

Note

You cannot upgrade the firmware for ATA disks.

When Data ONTAP takes disks offline

Data ONTAP temporarily stops I/O activity to a disk and takes a disk offline when:

- You update disk firmware
- ATA disks take a long time to recover from a bad media patch

While the disk is offline, Data ONTAP reads from other disks within the RAID group while writes are logged. The off-lined disk is brought back online after re-synchronization is complete. This process generally takes a few minutes and incurs a negligible performance impact. For ATA disks, this reduces the probability of forcing disk failures due to bad media patches or transient errors because disk-offline provides a software-based mechanism for fault-isolating drives and for performing out-of-band error recovery.

The disk offline feature is only supported for spares and data disks within RAID-DP and mirrored-RAID4 aggregates. A disk can be off-lined only if its containing RAID group is in a normal state and the plex or aggregate is not offline.

You view the status of disks with the `aggr status -r` or `aggr status -s` commands, as shown in the following examples. You can see what disks are offline with either option.

Note

For backward compatibility, you can also use the `vol status -r` or `vol status -s` commands.

Example 1:

```
toaster> aggr status -r aggrA
Aggregate aggrA (online, raid4-dp degraded) (block checksums)
Plex /aggrA/plex0 (online, normal, active)
RAID group /aggrA/plex0/rg0 (degraded)
RAID Disk Device HA SHELF BAY CHAN Pool Type RPM Used (MB/blks)
Phys (MB/blks)
--------- ------ ------------- ---- ---- ---- ----- --------------
--------------
parity 8a.20 8a 1 4 FC:A - FCAL 10000 1024/2097152 1191/2439568
data 6a.36 6a 2 4 FC:A - FCAL 10000 1024/2097152 1191/2439568
data 6a.19 6a 1 3 FC:A - FCAL 10000 1024/2097152 1191/2439568
data 8a.23 8a 1 7 FC:A - FCAL 10000 1024/2097152 1191/2439568
(offline)
```

Example 2:

```
toaster> aggr status -s
Spare disks
RAID Disk Device HA SHELF BAY CHAN Pool Type RPM Used (MB/blks)
Phys (MB/blks)
--------- ------ ------------- ---- ---- ---- ----- --------------
--------------
Spare disks for block or zoned checksum traditional volumes or
aggregates
spare 8a.24 8a 1 8 FC:A - FCAL 10000 1024/2097152 1191/2439568
spare 8a.25 8a 1 9 FC:A - FCAL 10000 1024/2097152 1191/2439568
spare 8a.26 8a 1 10 FC:A - FCAL 10000 1024/2097152 1191/2439568
(offline)
spare 8a.27 8a 1 11 FC:A - FCAL 10000 1024/2097152 1191/2439568
spare 8a.28 8a 1 12 FC:A - FCAL 10000 1024/2097152 1191/2439568
```

Disk access methods

About disk access methods

Several disk access methods are supported on NetApp appliances. This section discuses the following topics:

- "Multipath I/O for Fibre Channel disks" on page 49
- "Clusters" on page 54
- "Combined filer head disk shelf access" on page 55

Multipath I/O for Fibre Channel disks

Understanding Multipath I/O

The Multipath I/O feature for Fibre Channel disks enables you to create two paths, a primary path and a secondary path, from a single system to a disk loop. You can use this feature with or without SyncMirror.

Although it is not necessary to have a dual-port disk adapter to set up Multipath I/O, NetApp recommends you use two dual-port adapters to connect to two disk shelf loops, thus preventing either adapter from being the single point of failure. In addition, using dual-port adapters conserves Peripheral Component Interconnect (PCI) slots.

If your environment requires additional fault tolerance, you can use Multipath I/O with SyncMirror and configure it with four separate adapters, connecting one path from each adapter to one channel of a disk shelf. With this configuration, not only is each path supported by a separate adapter, but each adapter is on a separate bus. If there is a bus failure, or an adapter failure, only one path is lost.

Advantages of Multipath I/O

By providing redundant paths to the same disk on a single filer, the Multipath I/O feature offers the following advantages:

- Overall reliability and uptime of the storage subsystem of the filer is increased.
- Disk availability is higher.
- Bandwidth is increased (each loop provides an additional 200 MB/second of bandwidth).
- Storage subsystem hardware can be maintained with no downtime.

  When a primary host adapter is brought down, all traffic moves from that host adapter to the secondary host adapter. As a result, you can perform maintenance tasks, such as replacing a malfunctioning Loop Resiliency Circuit (LRC) module or cables connecting that host adapter to the disk shelves, without affecting the storage subsystem service.

Requirements to enable Multipath I/O on the filer

The Multipath I/O feature is enabled automatically, subject to the following restrictions:

- Only the following platforms support Multipath I/O:
  - F800 series
  - FAS900 series

  Note

  None of the NearStore appliance platforms (R100, R150, or R200 series) support Multipath I/O.

- Only the following host adapters support Multipath I/O:
  - QLOGIC 2200 (P/N X2040B)
  - QLOGIC 2212 (X2044A, 2044B)
  - QLOGIC 2342 (X2050A)
  - LSI 929X (X2050B)

  Note

  Although the 2200 and 2212 host adapters can co-exist with older (2100 and 2000) adapters on a filer, Multipath I/O is not supported on the older models.

  To determine the slot number where a host adapter can be installed in your filer, see the *System Configuration Guide* at the NOW site (http://now.netapp.com/).

- FC7 and FC8 disk shelves do not support Multipath I/O.
- FC9 must have two LRC modules to support Multipath I/O.
- DS14 and DS14mk2 FC disk shelves must have either two LRC modules or two Embedded Switch Hub (ESH) modules to support Multipath I/O.
- Older 9-GB disks (ST19171FC) and older 18-GB disks (ST118202FC) do not support Multipath I/O.
- Filers configured in clusters (that are not Fabric MetroClusters) do not support Multipath I/O.

  Multipath I/O setup and clustering setup both require the A and B ports of the disk shelves. Therefore, it is not possible to have both features enabled simultaneously.
- Hardware connections must be set up for Multipath I/O as specified in the corresponding *Fibre Channel StorageShelf* guide.

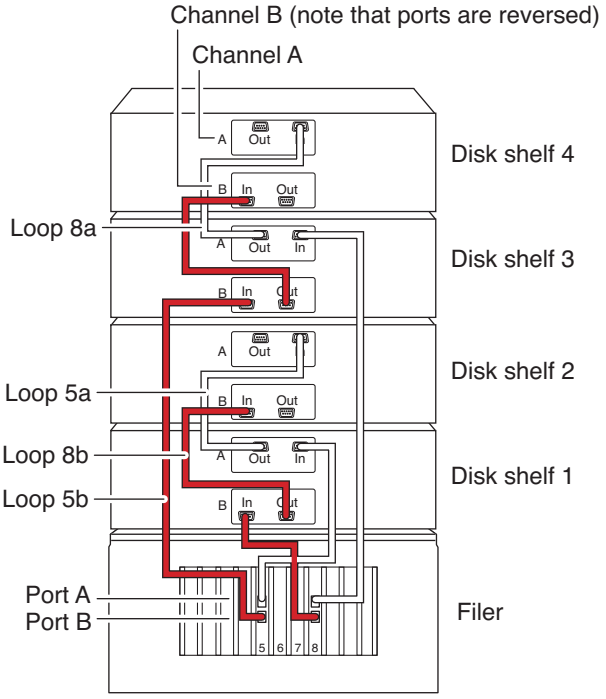
Supported configurations

Multipath I/O supports the following configuration:

- Multipath I/O without SyncMirror
- Multipath I/O with SyncMirror, not using software-based disk ownership
- Multipath I/O with SyncMirror, using software-based disk ownership
- Multipath I/O with SyncMirror, using four separate adapters

Multipath I/O without SyncMirror: NetApp recommends that you wire a single system for Multipath I/O without SyncMirror by connecting a primary path from one adapter to one disk loop and a secondary path from another adapter to that disk loop, as shown in the following illustration.

- The first loop is configured as follows:
  - Primary path: from system port 5a to disk shelves 1 and 2, A channels
  - Secondary path: from system port 8b to disk shelves 1 and 2, B channels
- The second loop is configured as follows:
  - Primary path: from system port 8a to disk shelves 3 and 4, A channels
  - Secondary path: from system port 5b to disk shelves 3 and 4, B channels

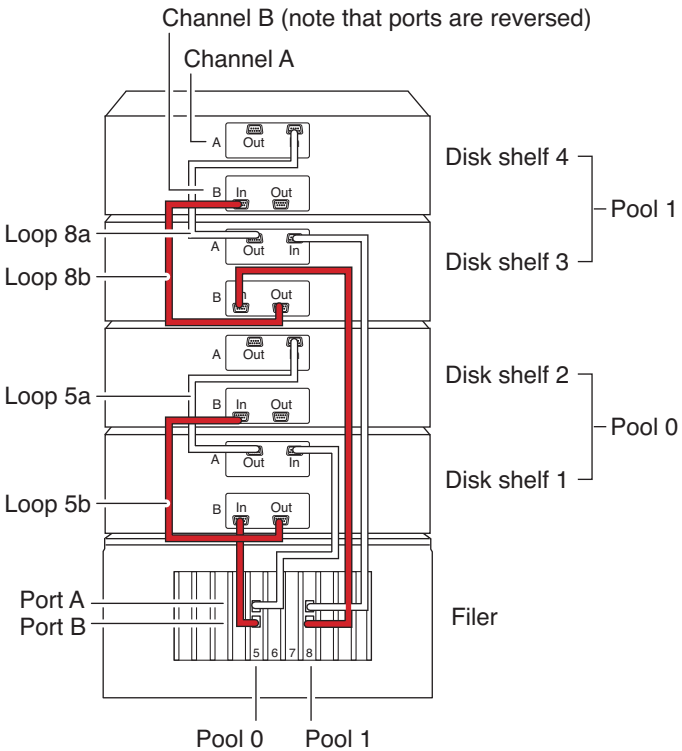


Multipath I/O without Syncmirror

Multipath I/O with SyncMirror not using software-based disk ownership: If your system does not support software-based disk ownership, you need to know which slots the adapters are in because that is what pool ownership is determined by. For example, with the FAS900 series, slots 1 through 7 own Pool0, and slots 8 through 11 own Pool1.

NetApp recommends that you configure the system to have a primary path and a secondary path connected from one adapter to the first disk loop and a primary and a secondary path from the other adapter to the second disk loop, as shown in the following illustration.

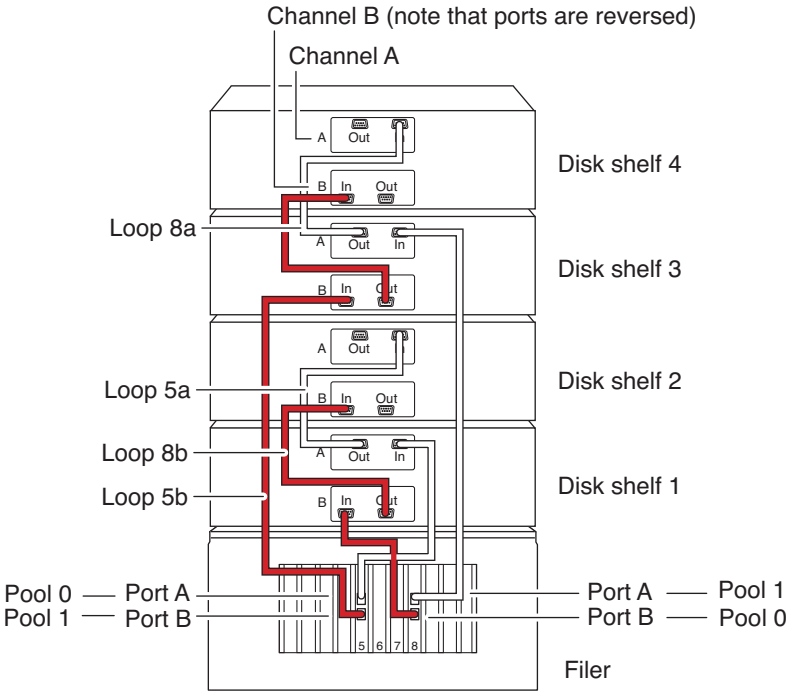
- The first loop is configured as follows:
  - Primary path: from system port 5a to disk shelves 1 and 2, A channels
  - Secondary path: from system port 5b to disk shelves 1 and 2, B channels
- The second loop is configured as follows:
  - Primary path: from system port 8a to disk shelves 3 and 4, A channels
  - Secondary path: from system port 8b to disk shelves 3 and 4, B channels



Multipath I/O with SyncMirror without software-based disk ownership

Multipath I/O with SyncMirror using software-based disk ownership: If your system supports software-based disk ownership, NetApp recommends that you configure the system to have a primary path and a secondary path from two different adapters to the first disk loop and a primary and a secondary path from the two adapters to the second disk loop, as shown in the following illustration.

- The first loop is configured as follows:
  - Primary path: from system port 5a to disk shelves 1 and 2, A channels
  - Secondary path: from system port 8b to disk shelves 1 and 2, B channels
    You can configure this as Pool0.
- The second loop is configured as follows:
  - Primary path: from system port 8a to disk shelves 3 and 4, A channels
  - Secondary path: from system port 5b to disk shelves 3 and 4, B channels
    You can configure this as Pool1.



Multipath I/O with Syncmirror with software-based disk ownership

Clusters

About clusters

NetApp clusters are two appliances in a partner relationship where each appliance can access the other's disk shelves as secondary owners. Each partner maintains two Fibre Channel Arbitrated Loops (or loops): a primary loop for a path to its own disks, and a secondary path to its partner's disk. The primary loop, loop A, is created by connecting the A ports of one or more disk shelves to the filer's disk adapter card, and the secondary loop, loop B, is created by connecting the B ports of one or more disk shelves to the filer's disk adapter card.

If one of the clustered appliances fails, its partner can start an emulated appliance that takes over serving the failed partner's disk shelves, providing uninterrupted access to its partner's disks as well as its own disks. For more information on installing clusters, see the *F800 and FAS900 Series Cluster Guide*; for information about managing clusters, see the *System Administration Guide*.

Combined filer head disk shelf access

About combined filer head and disk shelf appliances

Some NetApp appliances combine one or two filer heads and a disk shelf into a single unit. For example, the FAS270c consists of two clustered filer heads that share control of a single shelf of fourteen disks.

Primary clustered filer head ownership of each disk on the shelf is determined by software-based disk ownership information stored on each individual disk, not by A loop and B loop attachments. You use software-based disk ownership commands to assign each disk to the FAS270 filer heads, or any system with a SnapMover license.

For more information on software-based disk ownership assignment, see "Software-based disk ownership" on page 71.

Viewing information about all disks

To view information about all the disks connected to your filer, complete the following step.

Step	Action

1	Enter one of the following commands:
	`storage show disk`
	or
	`storage show disk -a`
	Result: Both commands display information about all disks in different formats. Without the `-a` option, the following information is displayed for all disks: Disk ID, shelf and bay location, serial number, vendor, model, and revision level.
	With the `-a` option, the following information is displayed for all disks: Disk ID, shelf and bay location, serial number, vendor, model, revision level, RPM (disk speed), WWN, Down rev (yes or no), primary port, power-on hours, blocks read, blocks written, time interval, Glist entries, scrub last done, and scrub count. The `-a` option displays information in a report form that is easily interpreted by scripts.

Example: The following example shows information about all the disks connected to the filer *toaster*:

```
toaster> storage show disk

DISK  SHELF  BAY    SERIAL     VENDOR   MODEL       REV
----  -----  ---    --------   ------   ---------   ----
7.6     0     6     LA774453   SEAGATE  ST19171FC   FB59
7.5     0     5     LA694863   SEAGATE  ST19171FC   FB59
7.4     0     4     LA781085   SEAGATE  ST19171FC   FB59
7.3     0     3     LA773189   SEAGATE  ST19171FC   FB59
7.14    1     6     LA869459   SEAGATE  ST19171FC   FB59
7.13    1     5     LA781479   SEAGATE  ST19171FC   FB59
7.12    1     4     LA772259   SEAGATE  ST19171FC   FB59
7.11    1     3     LA783073   SEAGATE  ST19171FC   FB59
7.10    1     2     LA700702   SEAGATE  ST19171FC   FB59
7.9     1     1     LA786084   SEAGATE  ST19171FC   FB59
7.8     1     0     LA761801   SEAGATE  ST19171FC   FB59
7.2     0     2     LA708093   SEAGATE  ST19171FC   FB59
7.1     0     1     LA773443   SEAGATE  ST19171FC   FB59
7.0     0     0     LA780611   SEAGATE  ST19171FC   FB59
```

Viewing the primary and secondary paths to the disks

To view the primary and secondary paths to all the disks connected to your filer, complete the following step.

Step	Action
1	Enter the following command: **`storage show disk -p`**

Note

The disk addresses shown for the primary and secondary paths to a disk are aliases of each other.

Example: In the following example, adapters are installed in the PCI expansion slot 7 and slot 8 of the filer tpubs-cf1. The slot 8 adapter is connected to port A of disk shelf 0, and the slot 7 adapter is connected to port B of disk shelf 1.

The output of the `storage show disk -p` command shows the primary and secondary paths to all disks connected to the filer.

```
tpubs-cf1> storage show disk -p
PRIMARY PORT  SECONDARY  PORT  SHELF BAY
------- ----  ---------- ----  ----- ---
7b.0     B    8a.0        A      0    0
8a.1     A    7b.1        B      0    1
8a.2     A    7b.2        B      0    2
7b.3     B    8a.3        A      0    3
7b.4     B    8a.4        A      0    4
8a.5     A    7b.5        B      0    5
8a.6     A    7b.6        B      0    6
7b.8     B    8a.8        A      1    0
8a.9     A    7b.9        B      1    1
7b.10    B    8a.10       A      1    2
7b.11    B    8a.11       A      1    3
7b.12    B    8a.12       A      1    4
7b.13    B    8a.13       A      1    5
8a.14    A    7b.14       B      1    6
```

Viewing the number of disks, spares, and failed disks

To see how many disks are on a system, open FilerView and go to Filer > Show Status. The following information is displayed: total number of disks, the number of spares, and the number of disks that have failed.

Viewing disk information in FilerView

You can use FilerView to select the types of disks to view information about: all disks, spare disks, broken disks, zeroing disks, and reconstructing disks. Open FilerView and go to Storage > Disks > Manage, and select the type of disk from the pull-down list. The following information about disks is displayed: Disk ID, type (parity, data, dparity, spare, and partner), checksum type, shelf and bay location, channel, size, physical size, pool, and aggregate.

Making disks available

How Data ONTAP initially makes disks available

Disks are recognized by Data ONTAP during bootup or when they are inserted in a disk shelf, unless you have an appliance that requires you to assign disks to the appliance, such as an FAS270 appliance or a gFiler gateway.

Once Data ONTAP recognizes a disk, it is initialized as a spare disk and automatically put into a pool and designated as a spare disk until it is assigned to a RAID group. All spare disks are in pool 0 unless the SyncMirror software is enabled. If SyncMirror is enabled, all spare disks are divided into two pools, Pool0 and Pool1.

The disks remain spare disks until they are designated as data disks or as parity disks by you or by Data ONTAP. Parity disks provide a level of protection against data loss if a disk in a RAID group fails.

Data ONTAP monitors the disks to ensure they are operating properly. When Data ONTAP detects disk failures, it takes corrective action and informs you of this activity.

Disk addressing

You identify a disk by its address, which is listed in the Device column of the `sysconfig -d` output. You can also use the `storage show` or `storage show disk -p` command to view the disk addresses. Disk addresses are represented in the following format:

path_id.device_id

path_id is generally the slot number on the filer where the host adapter is attached. However, certain exceptions exist. For example, the *path_id* for the onboard Fibre Channel card is 0a, and the *path_id* for the dual channel adapter is represented by *slot_number* a|b.

device_id is a protocol-specific identifier for the attached disks. For SCSI, the *device_id* is an integer from 0 to 15. For Fibre Channel-Arbitrated Loop (FC-AL), the *device_id* is an integer from 0 to 126.

SCSI disk addresses: For a SCSI disk, *path_id* is the disk's SCSI adapter number and *device_id* is the disk's SCSI ID number. For example, disks attached to the disk adapter 7a are numbered 7a.0 through 7a.14.

Fibre Channel disk addresses: For a Fibre Channel disk, *path_id* is the disk's host adapter number and *device_id* is the disk's ID. To identify the disk ID, see the `fcstat device_map` command output on your filer.

Available space on new disks

How Data ONTAP creates space consistency among disks from different manufacturers

When you add a new disk, Data ONTAP reduces the amount of space on that disk available for user data by rounding down. This maintains compatibility across disks from various manufacturers. The available disk space listed by informational commands such as `sysconfig` is, therefore, less for each disk than its rated capacity. The available disk space on a disk is rounded down as shown in the following table.

Disk	Right-sized Capacity	Available blocks
FC/SCSI disks		
4-GB disks	4 GB	8,192,000
9-GB disks	8.6 GB	17,612,800
18-GB disks	17 GB	34,816,000
35-GB disks (block checksum disks)	34 GB	69,632,000
36-GB disks (zoned checksum disks)	34.5 GB	70,656,000
72-GB disks	68 GB	139,264,000
144-GB disks	136 GB	278,528,000
288-GB disks	272 GB	557,056,000
ATA/SATA disks		
160-GB disks (available on R100 systems)	136 GB	278,258,000
250-GB disks (available on R150. R200, and FAS900 systems)	212 GB	434,176,000
320-GB disks (available on R200 and FAS900 systems)	274 GB	561,971,200

Displaying free disk space

You use the `df` command to verify the amount of free disk space on the filer.

Disk space report discrepancies

The total amount of disk space shown in the `df` output is less than the sum of available space on all disks installed in an aggregate.

In the following example, the `df` command is issued on a traditional volume with three 72-GB disks installed, with RAID-DP enabled, and the following data is displayed:

```
df /vol/vol0

Filesystem
                    kbytes     used     avail capacity Mounted on
/vol/vol0/        67108864   382296  66726568       1%
/vol/vol0/

/vol/vol0/.snapshot 16777216  14740  16762476       0%
/vol/vol/.snapshot
```

When you add the numbers in the kbytes column, the sum is significantly less than the total disk space installed. The following behavior accounts for the discrepancy:

- The two parity disks, which are 72-GB disks in this example, are not reflected in the output of the `df` command.
- The filer reserves 10 percent of the total disk space for efficiency, which `df` does not count as part of the file system space.

Note

The second line of output indicates how much space is allocated to snapshots. Snapshot reserve, if activated, can also cause discrepancies in the disk space report. For more information, see the *Data Protection Online Backup and Recovery Guide*.

Adding disks to the system

About adding disks

When you add disks to an appliance, you need to insert them in a disk shelf according to the instructions in the disk shelf manufacturer's documentation or the disk shelf guide provided by NetApp. When the disks are installed, they become hot-swappable spare disks, which means they can be replaced while the filer and shelves remain powered on.

Once the disks are recognized by Data ONTAP, you can add the disks to an aggregate with the `aggr add` command. For backward compatibility, you can also use the `vol add` command to add disks to the aggregate that contains a traditional volume.

Reasons to add disks to an appliance

You add disks for the following reasons:

- You want to add storage capacity to the appliance
- You are running out of hot spare disks
- You are replacing a disk

Considerations when adding disks to an appliance

The number of disks at initial appliance configuration affects read and write performance. A greater number of disks means a greater number of independently seeking disk-drive heads reading data, which improves performance. Write performance can also benefit from more disks; however, the difference can be masked by the effect of nonvolatile RAM (NVRAM) and the manner in which WAFL manages write operations.

As more disks are configured, the performance increase levels off. Performance is affected more with each new disk you add until the striping across all the disks levels out. When the striping levels out, there is an increase in the number of operations per second and a reduced response time.

For overall improved performance, NetApp recommends adding enough disks for a complete RAID group. The default RAID group size is NetApp system-specific.

When you add disks to an appliance that is a target in a SAN environment, NetApp recommends that you perform a full reallocation scan. For more information, see the *Block Access Management Guide for FCP* or the *Block Access Management Guide for iSCSI*.

When to add disks

To meet current storage requirements, NetApp recommends adding disks before a file system is 80 percent to 90 percent full.

To meet future storage requirements, NetApp recommends adding disks before the applied load places stress on the existing array of disks, even though adding more disks at this time will not significantly improve the filer's current performance immediately.

Adding new disks

Before adding new disks to the appliance, be sure that the filer supports the type of disk you want to add. For the latest information on supported disk drives, see the *Data ONTAP Release Notes* and the *System Configuration Guide* on the NOW site (http://now.netapp.com/).

Note

NetApp recommends that you add disks of the same size and the same checksum type, preferably block checksum.

To add new disks to the appliance, complete the following steps.

Step	Action	
1	**If the disks are...**	**Then...**
	Native Fibre Channel disks in Fibre Channel-attached shelves, or ATA disks on Fibre Channel-attached shelves	Go to Step 2.
	Native SCSI disks or ATA disks in SCSI-attached shelves	Enter the following command, and go to Step 2: `disk swap` **Note** For information on replacing a disk in an R100 or R200 series system, see the specific hardware and service guide for your system.

Step	Action
2	Install one or more disks according to the hardware guide for your disk shelf or the specific hardware and service guide for your system. **Note** On FAS270 and FAS270c appliances or systems licensed for SnapMover, a disk ownership assignment might need to be carried out. See “Software-based disk ownership” on page 71. **Result:** The system displays a message confirming that one or more disks were installed and then waits 15 seconds as the disks are recognized. The system recognizes the disks as hot spare disks. **Note** If you add multiple disks, the system might require 25 to 40 seconds to bring the disks up to speed as it checks the device addresses on each adapter.
3	Verify that the disks were added by entering the following command: `aggr status -s` **Result:** The number of hot spare disks in the RAID Disk column under Spare Disks increases by the number of disks you installed.

Adding hot spare disks to a filer

You should periodically check the number of hot spares you have in your appliance and add disks if you have fewer than one per shelf. If there are no hot spare disks, you should add disks to disk shelves so they become available as hot spares. For more information, see “Hot spare disks” on page 109.

Displaying the number of hot spare disks

To ascertain how many hot spare disks you have on your filer, complete the following step.

Step	Action

1	Enter the following command:
	aggr status -s
	Result: If there are hot spare disks, a display like the following appears, with a line for each spare disk, grouped by checksum type:
	`Pool1 spare disks`
	`RAID Disk Device HA SHELF BAY CHAN Pool Type RPM Used (MB/blks) Phys(MB/blks)`
	`--------- ----- ------------ ---- -------------- ----------`
	`Spare disks for block or zoned checksum traditonal volumes or aggregates`
	`spare 9a.24 9a 1 8 FC:A 1 FCAL 10000 34000/69532000 34190/70022840`
	`spare 9a.29 9a 1 13 FC:A 1 FCAL 10000 34000/69532000 34190/70022840`
	`Pool0 spare disks (empty)`

Removing disks

Reasons to remove disks

You remove a disk for the following reasons:

- You want to replace the disk because
  - It is a failed disk. You cannot use this disk again.
  - It is a data disk that is producing excessive error messages, and is likely to fail. You cannot use this disk again.
- You want to reuse the disk. It is a hot spare disk in good working condition, but you want to use it elsewhere.

You cannot reduce the number of disks in an aggregate by removing data disks. The only way to reduce the number of data disks in an aggregate is to copy the data and transfer it to a new file system that has fewer data disks.

Removing a failed disk

To remove a failed disk, complete the following steps.

Note

In the following procedure, the NearStore R100 and R150 appliance platforms support disks on SCSI-attached shelves. The R200 appliance platform supports disks on Fibre Channel-attached shelves. See the *Data ONTAP Release Notes* for the latest disk support information.

Step	Action
1	Find the disk ID of the failed disk by entering the following command: `aggr status -f` **Result:** The ID of the failed disk is shown next to the word *failed*. The location of the disk is shown to the right of the disk ID, in the column HA SHELF BAY.

Step	Action	
2	**If the disk is a...**	**Then...**
	Fibre Channel disk or in a Fibre Channel-attached shelf	Go to Step 3.
	SCSI disk or in a SCSI-attached shelf	Enter the following command and go to Step 3: `disk swap`
3	Remove the disk from the disk shelf according to the disk shelf manufacturer's instructions.	

Cancelling a disk swap command

To cancel the swap operation and continue service, complete the following step.

Step	Action
1	Enter the following command: `disk unswap`

Removing a hot spare disk

To remove a hot spare disk, complete the following steps.

Note

In the following procedure, the NearStore R100 and R150 appliance platforms support disks on SCSI-attached shelves. See the *Data ONTAP Release Notes* for the latest disk support information.

Step	Action
1	Find the disk IDs of hot spare disks by entering the following command: `aggr status -s` **Result:** The names of the hot spare disks appear next to the word *spare*. The locations of the disks are shown to the right of the disk name.

Step	Action	
2	Enter the following command to spin down the disk: **`disk remove disk_name`** *disk_name* is the name of the disk you want to remove (from the output of Step 1).	
3	**If the disk is...**	**Then...**
	A Fibre Channel disk or in a Fibre Channel-attached shelf	Go to Step 4.
	A SCSI disk or in a SCSI-attached shelf	Enter the following command, and go to Step 4: **`disk swap`**
4	Wait for the disk to stop spinning. See the hardware guide for your disk shelf model for information about how to tell when a disk stops spinning.	
5	Remove the disk from the disk shelf, following the instructions in the hardware guide for your disk shelf model. **Result:** When replacing FC disks, there is no service interruption. When replacing SCSI and ATA disks, file service resumes 15 seconds after you remove the disk.	

Removing a data disk

To remove a data disk, complete the following steps.

Note

In the following procedure, the NearStore R100 and R150 appliance platforms support disks on SCSI-attached shelves. See the *Data ONTAP Release Notes* for the latest disk support information.

Step	Action
1	Find the disk name in the log messages that report disk errors by looking at the numbers that follow the word *Disk*.
2	Enter the following command: **`aggr status -r`**
3	Look at the Device column of the output of the `sysconfig -r` command. It shows the disk ID of each disk. The location of the disk appears to the right of the disk ID, in the column HA SHELF BAY.
4	Enter the following command to fail the disk: **`disk fail [-i]`** ***`disk_name`*** `-i` specifies to fail the disk immediately. *disk_name* is the disk name from the output in Step 1.

<table>
<tr><th>Step</th><th colspan="2">Action</th></tr>
<tr><td></td><td>If you...</td><td>Then...</td></tr>
<tr><td></td><td>Do not specify the <code>-i</code> option</td><td>Data ONTAP pre-fails the specified disk and attempts to create a replacement disk by copying the contents of the pre-failed disk to a spare disk.

This copy might take several hours, depending on the size of the disk and the load on the appliance.

Caution
You must wait for the disk copy to complete before going to the next step.

If the copy operation is successful, then the pre-failed disk is failed and the new replacement disk takes its place.</td></tr>
<tr><td></td><td>Specify the <code>-i</code> option or if the disk copy operation fails</td><td>The pre-failed disk fails and the system operates in degraded mode until the RAID system reconstructs a replacement disk.</td></tr>
<tr><td>5</td><td>If the disk is...</td><td>Then...</td></tr>
<tr><td></td><td>A Fibre Channel disk or in a Fibre Channel-attached shelf</td><td>Go to Step 6.</td></tr>
<tr><td></td><td>A SCSI disk or in a SCSI-attached shelf</td><td>Enter the following command, then go to Step 6:
<code>disk swap</code></td></tr>
<tr><td>6</td><td colspan="2">Remove the failed disk from the disk shelf, following the instructions in the hardware guide for your disk shelf model.

Result: File service resumes 15 seconds after you remove the disk.</td></tr>
</table>

Software-based disk ownership

About software-based disk ownership

Software-based disk ownership software assigns ownership of a disk to a specific filer head by writing software ownership information on the disk rather than by using the topology of the storage system's physical connections. Software-based disk ownership is implemented in systems where a disk shelf can be accessed by more than one filer head. Configurations that use software-based disk ownership include

- FAS270 systems
- Clusters configured for SnapMover vFiler™ migration. For more information, see the section on the SnapMover vFiler no copy migration feature in the *MultiStore Management Guide*.
- gFiler arrays. For more information, see the section on SnapMover in the *gFiler Gateway Series Software Setup, Installation, and Management Guide*.

FAS270 systems

The NetApp FAS270 and FAS270c systems consist of a single disk shelf of 14 disks and either one internal filer head (on the FAS270) or two clustered internal filer heads (on the FAS270c). By design, a disk located on this common disk shelf can, if the appliance has two filer heads, be assigned to the ownership of either filer head. The ownership of each disk is ascertained by an ownership record written on each disk.

NetApp delivers the FAS270 and FAS270c appliances with each disk preassigned to the single FAS270 internal filer head or preassigned to one of the two FAS270c filer heads.

However, if you add one or more disk shelves to an existing FAS270 or FAS270c appliance, you might have to assign ownership of the disks contained on those shelves.

Assigning disks

To assign disks that are currently labeled "not owned," complete the following steps.

Step	Action
1	Use the `disk show -n` command to view all disks that do not have assigned owners.

Step	Action
2	Use the following command to assign the disks that are labeled “not owned” to one of the filer heads. If you are assigning unowned disks to a non-local filer, you must identify the filer by using either the `-o ownername` or the `-s sysid` parameters, or both. **`disk assign {`*`disk_name`*` |all| -n count} [-p pool] [-o ownername] [-s sysid] [-c block|zoned] [-f]`** *disk_name* specifies the disk that you want to assign to the filer. `all` specifies all of the unowned disks are assigned to the filer. `-n` *`count`* specifies the number of unassigned disks to be assigned to the filer, as specified by `count`. `-p` *`pool`* specifies which SyncMirror pool the disks are assigned to. The value of pool is either 0 or 1. `-o` *`ownername`* specifies the filer that the disks are assigned to. `-s sysid` specifies the filer that the disks are assigned to. `-c` specifies the checksum type (either block or zoned) for a LUN in gFiler gateways. `-f` must be specified if the filer already owns the disk. **Example:** The following command assigns six disks on the FAS270c to the filer fh1: `fh1>` **`disk assign 0b.43 0b.41 0b.39 0b.37 0b.35 0b.33`** **Result:** The specified disks are assigned as disks to the filer on which the command was executed.
3	Use the `disk show -v` command to verify the disk assignments that you have just made.

After you have assigned disks, you can assign those disks to the aggregates on the filer that owns them, or leave them as spare disks on that filer.

Viewing disk ownership

To view the ownership of disks on the FAS270 or the FAS270c or systems licensed for SnapMover, complete the following step.

Step	Action
1	Enter the following command to display a list of all the disks. `fh1> **disk show -v**`

Sample output: The following sample output of the `disk show -v` command on an FAS270c shows disks 0b.16 through 0b.29 assigned in odd/even fashion to the internal cluster nodes (or filer heads) fh1and fh2. The 14 disks on the add-on disk shelf are still unassigned to either filer head.

```
fh1> disk show -v
 DISK      OWNER              POOL   SERIAL NUMBER
 --------- ---------------    -----  -------------
 0b.43     Not Owned          NONE   41229013
 0b.42     Not Owned          NONE   41229012
 0b.41     Not Owned          NONE   41229011
 0b.40     Not Owned          NONE   41229010
 0b.39     Not Owned          NONE   41229009
 0b.38     Not Owned          NONE   41229008
 0b.37     Not Owned          NONE   41229007
 0b.36     Not Owned          NONE   41229006
 0b.35     Not Owned          NONE   41229005
 0b.34     Not Owned          NONE   41229004
 0b.33     Not Owned          NONE   41229003
 0b.32     Not Owned          NONE   41229002
 0b.31     Not Owned          NONE   41229001
 0b.30     Not Owned          NONE   41229000
 0b.29     fh1 (84165672)     Pool0  41226818
 0b.28     fh2 (84165664)     Pool0  41221622
 0b.27     fh1 (84165672)     Pool0  41226333
 0b.26     fh2 (84165664)     Pool0  41225544
 0b.25     fh1 (84165672)     Pool0  41221700
 0b.24     fh2 (84165664)     Pool0  41224003
 0b.23     fh1 (84165672)     Pool0  41227932
 0b.22     fh2 (84165664)     Pool0  41224591
 0b.21     fh1 (84165672)     Pool0  41226623
 0b.20     fh2 (84165664)     Pool0  41221819
 0b.19     fh1 (84165672)     Pool0  41227336
 0b.18     fh2 (84165664)     Pool0  41225345
 0b.17     fh1 (84165672)     Pool0  41225446
 0b.16     fh2 (84165664)     Pool0  41201783
```

Additional `disk show` options are listed below.

`disk show` parameter	**Information displayed**
`disk show -a`	Shows all disks that have an owner
`disk show [-o` *`hostname`* `| -s` *`nvram_id`* `]`	Shows all disks owned by a specified host or its NVRAM ID (the NVRAM number shipped with the filer head)
`disk show -n`	Shows all the disks that do not have an owner
`disk show -v`	Shows all the visible disks

Modifying disk assignments

You can also use the `disk assign` command to modify the ownership of any disk assignment that you have made. For example, on the FAS270c appliance, you can reassign a disk from one filer head to the other. On either the FAS270 or FAS270c appliance, you can change an assigned disk back to “not owned” status.

Caution

NetApp recommends that disk assignments be modified only for spare disks. Disks that have already been assigned to an aggregate cannot be reassigned without endangering all the data and the structure of that entire aggregate.

To modify disk ownership assignments on the FAS270 or the FAS270c, or systems licensed for SnapMover, complete the following steps.

Step	**Action**
1	In the command-line interface of the FAS270 or FAS270c filer head, or systems licensed for SnapMover whose disks you want to reassign, use the `aggr status -r` command to view the spare disks, whose ownership can safely be changed.

Step	Action
2	Use the following command to modify assignment of the spare disks. **`disk assign {`*`disk1`* **`[`***`disk2`***`] [...]|-n`** *`num_disks`***`} -f {-o`** *`ownername`* **`| -s unowned	-s`** *`sysid`***`}`** *disk1* [*disk2*] [...] are the names of the spare disks whose ownership assignment you want to modify. `-n` *num_disks* specifies a number of disks, rather than a series of disk names, to assign ownership to. `-f` is a switch required for forcing the assignment of disks that have already been assigned ownership. `-o` *ownername* specifies the host name of the filer head to which you want to reassign the disks in question. `-s unowned` modifies the ownership assignment of the disks in question back to "not owned." `-s` *sysid* is the factory-assigned NVRAM number of the filer head to which you want to reassign the disks. **Example:** The following command unassigns four disks on the FAS270c from the filer fh1: `fh1>` **`disk assign 0b.30 0b.29 0b.28 0b.27 -s unowned`**
3	Use the `disk show -v` command to verify the disk assignment modifications that you have just made.

Re-using disks configured for software-based disk ownership

Re-using disks that are configured for software-based disk ownership

If you want to re-use disks from NetApp appliances that have been configured for software-based disk ownership, NetApp strongly recommends that you take precautions if you reinstall these disks in appliances that do not use software-based disk ownership. Disks with unerased software-based ownership information that are installed in an unbooted appliance that does not use software-based disk ownership will cause that appliance to fail on reboot. Precautions include:

- Erasing the software-based disk ownership information from a disk prior to removing it from its original system. See "Erasing software-based disk ownership prior to removing a disk" on page 77.
- Transferring the disks to the target system while that system is in operation. See "Automatically erasing disk ownership information" on page 77.
- If you accidentally cause a boot failure by installing software-assigned disks, undoing this mishap by running the `disk remove_ownership` command in maintenance mode. See "Undoing accidental conversion to software-based disk ownership" on page 78.

Erasing software-based disk ownership prior to removing a disk

If possible, NetApp recommends that you erase software-based disk ownership information on the target disks before removing them from their current appliance and prior to transferring them to another system.

To undo software-based disk ownership on a target disk prior to removing it, complete the following steps.

Step	Action
1	At the prompt of the appliance whose disks you want to transfer, enter the following command to list all the appliance disks and their RAID status. `aggr status -r` Note the names of the disks that you want to transfer. **Note** In most cases, (unless you plan to physically move an entire aggregate of disks to a new system) you should plan to transfer only disks listed as hot spare disks.
2	For each disk that you want to remove, enter the following command: `disk remove_ownership` ***`disk_name`*** *disk_name* is the name of the disk whose software-based ownership information you want to remove.
3	Enter the following command to confirm the removal of the disk ownership information from the specified disk. `disk show -v` **Result:** The specified disk and any other disk that is labeled “not owned” is ready to be moved to other systems.
4	Remove the specified disk from its original system and install it into its target system.

Automatically erasing disk ownership information

If you physically transfer disks from an appliance that uses software-based disk ownership to a running appliance that does not, you can do so without using the `disk remove_ownership` command if that appliance is running Data ONTAP 6.5.1 or higher.

Step	Action	
1	Do not shut down the target appliance.	
2	On the target appliance, enter the following command to confirm the version of Data ONTAP on the target appliance. `version`	
3	**If**	**Then**
	The Data ONTAP version on the target appliance is 6.5.1 or later	Go to Step 4.
	The Data ONTAP version on the target appliance is earlier than 6.5.1	Do not continue this procedure; instead, erase the software-based disk ownership information on the source appliance, as described in "Erasing software-based disk ownership prior to removing a disk" on page 77.
4	Remove the disks from their original appliance and physically install them in the running target appliance. If Data ONTAP 6.5.1 or later is installed, the running target appliance will automatically erase any existing software-based disk ownership information on the transferred disks.	
5	On the target appliance, use the `aggr status -r` command to verify that the disks you have added are successfully installed.	

Undoing accidental conversion to software-based disk ownership

If you transfer disks from a system configured for software-based disk ownership (such as the FAS270 appliance, or a cluster enabled for SnapMover vFiler™ migration) to another system that does not use software-based disk ownership, you might accidentally mis-configure that target system as a result of the following circumstances.

- You neglect to remove software-based disk ownership information from the target disks before you remove them from their original system.

- You add the disks to a target system that does not use software-based disk ownership while the target system is off.
- The target system is upgraded to Data ONTAP 6.5.1 or later.

Under these circumstances, if you reboot the target system in normal mode, the remaining disk ownership information causes the target system to convert to a mis-configured software-based disk ownership setup. It will fail reboot.

To undo this accidental conversion to software-based disk ownership, complete the following steps.

Step	Action
1	Turn on or reboot the target system. When prompted to do so, press Ctrl-C to display the boot menu.
2	Enter the choice for booting in maintenance mode.
3	In maintenance mode, enter the following command: `disk remove_ownership all` The software-based disk ownership information is erased from all disks that have them.
4	Halt the system to exit maintenance mode by entering the following command: `halt`
5	Reboot the target system. The system will reboot in normal mode with software-based disk ownership disabled.

Disk sanitization

About disk sanitization

Disk sanitization is the process of physically obliterating data by overwriting disks with specified byte patterns or random data in a manner that prevents recovery of the original data by any known recovery methods. You sanitize disks if you want to ensure that data currently on those disks is physically unrecoverable. For example, you might have some disks that you intend to remove from one appliance and you want to re-use those disks in another appliance or simply dispose of the disks. In either case, you want to ensure no one can retrieve any data from those disks.

The Data ONTAP `disk sanitize` command enables you to carry out disk sanitization by using three successive default or user-specified byte overwrite patterns for up to seven cycles per operation. You can start, stop, and display the status of the disk sanitization process, which runs in the background. Depending on the capacity of the disk and the number of patterns and cycles specified, this process can take several hours to complete. When the process has completed, the disk is in a sanitized state. You can return a sanitized disk to the spare disk pool with the `disk sanitize release` command.

What this section covers

The section covers the following topics:

- "Disk sanitization limitations" on page 80
- "Licensing disk sanitization" on page 81
- "Sanitizing disks" on page 82
- "Stopping disk sanitization" on page 85
- "Selectively sanitizing data" on page 85
- "Reading disk sanitization log files" on page 90

Disk sanitization limitations

The disk sanitization operations described in this section are subject to the following limitations:

- The operations are not supported on gFiler appliances.
- The operations are not supported in takeover mode on clustered systems, or on disabled cluster systems.
- Disk sanitization cannot be carried out on disks that were failed due to readability or writability problems.

- Disk sanitization cannot be carried out on disks that belong to an SEC 17a-4-compliant SnapLock volume until the expiration periods on all files have expired, that is, all of the files have reached their retention dates.
- The formatting phase of the disk sanitization process is skipped on ATA drives.
- NetApp does not support the disk sanitization process on older disks.

  To determine if disk sanitization is supported on a specified disk, run the `storage show disk` command. If the vendor for the disk in question is listed as NETAPP, disk sanitization is supported.

Licensing disk sanitization

Before you can use the disk sanitization feature, you must install the disk sanitization license.

Caution

Once installed on a NetApp system, the license for disk sanitization is permanent.

Caution

The disk sanitization license prohibits the following admin command from being used on the NetApp system:

- `dd` (to copy blocks of data)

Caution

The disk sanitization license prohibits the following diagnostic commands from being used on the NetApp system:

- `dumpblock` (to print dumps of disk blocks)
- `setflag wafl_metadata_visible` (to allow access to internal WAFL files)

To install the disk sanitization license, complete the following step:

Step	Action
1	Enter the following command: **`license add`** ***`disk_sanitize_code`*** *disk_sanitize_code* is the disk sanitization license code that NetApp provides.

Sanitizing disks

You can sanitize any disk that has *spare* status. This includes disks that exist on the appliance as spare disks after the aggregate that they belong to has been destroyed. It also includes disks that were removed from the spare disk pool by the `disk remove` command but have been returned to spare status after an appliance reboot.

To sanitize a disk or a set of disks on an appliance, complete the following steps.

Step	Action
1	Print a list of all disks assigned to RAID groups, failed, or existing as spares, by entering the following command. **`sysconfig -r`** Do this to verify that the disk or disks that you want to sanitize do not belong to any existing RAID group in any existing aggregate.
2	Enter the following command to sanitize the specified disk or disks of all existing data. **`disk sanitize start [-p pattern1|-r [-p pattern2|-r] [-p pattern|-r]]] [-c cycle_count] disk_list`** `-p` *pattern1* `-p` *pattern2* `-p` *pattern3* specifies a cycle of one to three user-defined hex byte overwrite patterns that can be applied in succession to the disks being sanitized. The default hex pattern specification is `-p 0x55 -p 0xAA -p 0x3c`. `-r` replaces a patterned overwrite with a random overwrite for any or all of the cycles, for example: `-p 0x55 -p 0xAA -r` `-c` *cycle_count* specifies the number of cycles for applying the specified overwrite patterns. The default value is one cycle. The maximum value is seven cycles. **Note** To be in compliance with United States Department of Defense and Department of Energy security requirements, you must set *cycle_count* to six cycles per operation. *disk_list* specifies a space-separated list of spare disks to be sanitized.

Step	Action

	Example: The following command applies the default three disk sanitization overwrite patterns for one cycle (for a total of 3 overwrites) to the specified disks, 7.6, 7.7, and 7.8.
	`disk sanitize start 7.6 7.7 7.8`
	If you set *cycle_count* to 6, this example would result in three disk sanitization overwrite patterns for six cycles (for a total of 18 overwrites) to the specified disks.
	Result: The specified disks are sanitized, put into the pool of broken disks, and marked as *sanitized*. A list of all the sanitized disks is stored in the appliance's `/etc` directory.
	Note
	If you need to abort the sanitization operation, enter
	`disk sanitize abort [disk_list]`
	Caution
	Do not turn off the appliance, disrupt the disk loop, or remove target disks during the sanitization process. If the sanitization process is disrupted, the target disks that are in the formatting stage of disk sanitization will require reformatting before their sanitization can be completed. See "If formatting is interrupted" on page 85.

3	To check the status of the disk sanitization process, enter the following command: `disk sanitize status [disk_list]`

Step	Action
4	To release sanitized disks from the pool of broken disks for reuse as spare disks, enter the following command: **`disk sanitize release` *`disk_list`*** **Caution** The `disk sanitize release` command removes the *sanitized* label from the affected disks and returns them to spare state. Rebooting the filer or removing the disk also removes the sanitized label from any sanitized disks and returns them to spare state. **Verification:** To list all disks on the appliance and verify the release of the sanitized disks into the pool of spares, enter **`sysconfig -r.`**

Process description: After you enter the `disk sanitize start` command, Data ONTAP begins the sanitization process on each of the specified disks. The process consists of a disk format operation, followed by the specified overwrite patterns repeated for the specified number of cycles.

Note
The formatting phase of the disk sanitization process is skipped on ATA disks.

The time to complete the sanitization process for each disk depends on the size of the disk, the number of patterns specified, and the number of cycles specified.

For example, the following command invokes one format overwrite pass and 18 pattern overwrite passes of disk 7.3.

```
disk sanitize start -p 0x55 -p 0xAA -p 0x37 -c 6 7.3
```

- If disk 7.3 is 36 GB and each formatting or pattern overwrite pass on it takes 15 minutes, then the total sanitization time is 19 passes times 15 minutes, or 285 minutes (4.75 hours).
- If disk 7.3 is 73 GB and each formatting or pattern overwrite pass on it takes 30 minutes, then total sanitization time is 19 passes times 30 minutes, or 570 minutes (9.5 hours).

If disk sanitization is interrupted: If the sanitization process is interrupted by power failure, filer panic, or a user-invoked `disk sanitize abort` command, the `disk sanitize` command must be re-invoked and the process repeated from the beginning in order for the sanitization to take place.

If formatting is interrupted: If the formatting phase of disk sanitization is interrupted, Data ONTAP attempts to reformat any disks that were corrupted by an interruption of the formatting. After a system reboot and once every hour, Data ONTAP checks for any sanitization target disk that did not complete the formatting phase of its sanitization. If such a disk is found, Data ONTAP attempts to reformat that disk, and writes a message to the console informing you that a corrupted disk has been found and will be reformatted. After the disk is reformatted, it is returned to the hot spare pool. You can then rerun the `disk sanitize` command on that disk.

Stopping disk sanitization

You can use the `disk sanitize abort` command to stop an ongoing sanitization process on one or more specified disks. If you use the `disk sanitize abort` command, the specified disk or disks are returned to spare state and the sanitized label is removed.

Step	Action
1	Enter the following command: **`disk sanitize abort`** ***`disklist`*** Data ONTAP displays the message “Sanitization abort initiated.” If the specified disks are undergoing the disk formatting phase of sanitization, the abort will not occur until the disk formatting is complete. Once the process is stopped, Data ONTAP displays the message “Sanitization aborted for *diskname*.”

Selectively sanitizing data

Selective data sanitization consists of physically obliterating data in specified blocks while preserving all other data located on the affected aggregate for continued user access.

Summary of the selective sanitization process: Because data for any one file on an appliance is physically stored on any number of data disks in the aggregate containing that data, and because the physical location of data within an aggregate can change, sanitization of selected data, such as files or directories, requires that you sanitize every disk in the aggregate where the data is located (after first migrating the aggregate data that you do *not* want to sanitize to disks on another aggregate). To selectively sanitize data contained in an aggregate, you must carry out three general tasks.

1. Delete the selected files or directories (and any aggregate snapshots that contain those files or directories) from the aggregate that contains them.
2. Migrate the remaining data (the data that you want to preserve) in the affected aggregate to a new set of disks in a destination aggregate on the same appliance using `ndmpcopy` command.
3. Destroy the original aggregate and sanitize all the disks that were RAID group members in that aggregate.

Requirements for selective sanitization: Successful completion of this process requires the following conditions:

- You must install a disk sanitization license on your appliance.
- You must have enough storage space on your appliance to create an additional destination aggregate to which you can migrate the data that you want to preserve from the original aggregate. This destination aggregate must have a storage capacity at least as large as that of the original aggregate.
- You must use the `ndmpcopy` command to migrate data in the affected aggregate to a new set of disks in a destination aggregate on the same appliance. For information about the `ndmpcopy` command, see the *Data Protection Online Backup and Recovery Guide*.

Aggregate size and selective sanitization: Because sanitization of any unit of data in an aggregate still requires you to carry out data migration and disk sanitization processes on that entire aggregate, NetApp recommends that you use small aggregates to store data that requires sanitization. Use of small aggregates for storage of data requiring sanitization minimizes the time, disk space, and bandwidth that sanitization will requires.

Backup and data sanitization: Absolute sanitization of data means physical sanitization of *all* instances of aggregates containing sensitive data; it is therefore advisable to maintain your sensitive data in aggregates that are *not* regularly backed up to aggregates that also back up large amounts of nonsensitive data.

Procedure for selective sanitization: To carry out selective sanitization of data within an aggregate or a traditional volume, complete the following steps.

Step	Action
1	From a Windows or UNIX client, delete the directories or files whose data you want to selectively sanitize from the active file system. Use the appropriate Windows or UNIX command, such as **`rm -rf /nixdir/nixfile.doc`**

Step	Action
2	From the NetApp appliance, enter the following commands to delete all snapshots of the aggregates and volumes (both traditional and FlexVol volumes) that contain the files or directories that you just deleted. ◆ To delete all snapshots associated with the aggregate, enter the following command: **`snap delete -a aggr_name -A`** *aggr_name* is the aggregate that contains the files or directories that you just deleted. For example: **`snap delete -a nixsrcaggr -A`** ◆ To delete all snapshots associated with the volume, enter the following command: **`snap delete -a vol_name -V`** *vol_name* is the traditional volume or FlexVol that contains the files or directories that you just deleted. For example: **`snap delete -a nixsrcvol -V`** ◆ To delete a specific snapshot for either an aggregate or a volume, enter one of the following commands: **`snap delete aggr_name snapshot_name -A`** **`snap delete vol_name snapshot_name -V`** Examples: **`snap delete nixsrcaggr nightly0 A`** **`snap delete nixsrcvol nightly0 -V`**
3	Enter the following command to determine the size of the aggregate from which you deleted data: **`aggr status aggr_name -b`** For backward compatibility, you can also use the following command for traditional volumes. **`vol status vol_name -b`** **Example:** **`aggr status nixsrcaggr -b`** Calculate the aggregate size in bytes by multiplying the bytes per block (block size) by the blocks per aggregate (aggregate size).

Step	Action
4	Enter the following command to create an aggregate to which you will migrate undeleted data. This aggregate must be of equal or greater storage capacity than the aggregate from which you just deleted file, directory, or snapshot data: **`aggr create dest_aggr ndisks`** For backward compatibility with traditional volumes, you can also enter: **`vol create dest_vol disklist`** **Example:** **`aggr create nixdestaggr 8@72G`** **Note** The purpose of this new aggregate is to provide a migration destination that is absolutely free of the data that you want to sanitize.
5	Enter the following command to copy the data you want to preserve to the destination aggregate from the source aggregate you want to sanitize: **`ndmpcopy src_aggr dest_aggr`** *src_aggr* is the source aggregate. *dest_aggr* is the destination aggregate. **Caution** Be sure that you have deleted the files or directories that you want to sanitize (and any affected snapshots) from the source aggregate *before* you run the `ndmpcopy` command. **Example:** **`ndmpcopy nixsrcvol nixdestvol`**

Step	Action
6	Record the disks currently in the source aggregate. (After that aggregate is destroyed, you will sanitize these disks.) To list the disks in the source aggregate, enter the following command: **`aggr status src_aggr -r`** **Example:** **`aggr status nixsrcaggr -r`** The disks that you are going to sanitize are listed in the Device column of the `aggr status -r` output.
7	In maintenance mode, enter the following command to take the source aggregate offline: **`aggr offline src_aggr`** **Example:** **`aggr offline nixsrcaggr`**
8	Enter the following command to destroy the source aggregate: **`aggr destroy src_aggr`** **Example:** **`aggr destroy nixsrcaggr`**
9	Enter the following command to rename the destination aggregate, giving it the name of the source aggregate that you just destroyed: **`aggr rename dest_aggr src_aggr`** **Example:** **`aggr rename nixdestaggr nixsrcaggr`**
10	Reestablish your CIFS or NFS services: ◆ If the original volume supported CIFS services, restart the CIFS services on the volumes in the destination aggregate after migration is complete. ◆ If the original volume supported NFS services, enter the following command: **`exportfs -a`** **Result:** Users who were accessing files in the original volume will continue to access those files in the renamed destination volume with no remapping of their connections required.

Step	Action
11	Use the `disk sanitize` command to sanitize the disks that used to belong to the source aggregate. Follow the procedure described in “Sanitizing disks” on page 82.

Reading disk sanitization log files

The disk sanitization process outputs two types of log files.

- One file, /etc/sanitized_disks, lists all the drives that have been sanitized.
- For each disk being sanitized, a file is created where the progress information will be written.

Listing the sanitized disks: The /etc/sanitized_disks file contains the serial numbers of all drives that have been successfully sanitized. For every invocation of the `disk sanitize start` command, the serial numbers of the newly sanitized disks are appended to the file.

The /etc/sanitized_disks file shows output similar to the following:

```
admin1> rdfile /etc/sanitized_disks
Tue Jun 24 02:54:11 Disk 8a.44 [S/N 3FP0RFAZ00002218446B]
sanitized.
Tue Jun 24 02:54:15 Disk 8a.43 [S/N 3FP20XX400007313LSA8]
sanitized.
Tue Jun 24 02:54:20 Disk 8a.45 [S/N 3FP0RJMR0000221844GP]
sanitized.
Tue Jun 24 03:22:41 Disk 8a.32 [S/N 43208987] sanitized.
```

Reviewing the disk sanitization progress: A progress file is created for each drive sanitized and the results are consolidated to the /etc/sanitization.log file every 15 minutes during the sanitization operation. Entries in the log resemble the following:

```
Tue Jun 24 02:40:10 Disk sanitization initiated on drive 8a.43 [S/N
3FP20XX400007313LSA8]
Tue Jun 24 02:40:10 Disk sanitization initiated on drive 8a.44 [S/N
3FP0RFAZ00002218446B]
Tue Jun 24 02:40:10 Disk sanitization initiated on drive 8a.45 [S/N
3FP0RJMR0000221844GP]
Tue Jun 24 02:53:55 Disk 8a.44 [S/N 3FP0RFAZ00002218446B] format
completed in 00:13:45.
Tue Jun 24 02:53:59 Disk 8a.43 [S/N 3FP20XX400007313LSA8] format
completed in 00:13:49.
Tue Jun 24 02:54:04 Disk 8a.45 [S/N 3FP0RJMR0000221844GP] format
```

```
completed in 00:13:54.
Tue Jun 24 02:54:11 Disk 8a.44 [S/N 3FP0RFAZ00002218446B] cycle 1
pattern write of 0x47 completed in 00:00:16.
Tue Jun 24 02:54:11 Disk sanitization on drive 8a.44 [S/N
3FP0RFAZ00002218446B] completed.
Tue Jun 24 02:54:15 Disk 8a.43 [S/N 3FP20XX400007313LSA8] cycle 1
pattern write of 0x47 completed in 00:00:16.
Tue Jun 24 02:54:15 Disk sanitization on drive 8a.43 [S/N
3FP20XX400007313LSA8] completed.
Tue Jun 24 02:54:20 Disk 8a.45 [S/N 3FP0RJMR0000221844GP] cycle 1
pattern write of 0x47 completed in 00:00:16.
Tue Jun 24 02:54:20 Disk sanitization on drive 8a.45 [S/N
3FP0RJMR0000221844GP] completed.
Tue Jun 24 02:58:42 Disk sanitization initiated on drive 8a.43 [S/N
3FP20XX400007313LSA8]
Tue Jun 24 03:00:09 Disk sanitization initiated on drive 8a.32 [S/N
43208987]
Tue Jun 24 03:11:25 Disk 8a.32 [S/N 43208987] cycle 1 pattern write
of 0x47 completed in 00:11:16.
Tue Jun 24 03:12:32 Disk 8a.43 [S/N 3FP20XX400007313LSA8]
sanitization aborted by user.
Tue Jun 24 03:22:41 Disk 8a.32 [S/N 43208987] cycle 2 pattern write
of 0x47 completed in 00:11:16.
Tue Jun 24 03:22:41 Disk sanitization on drive 8a.32 [S/N 43208987]
completed.
```

Disk and storage subsystem management

Command for viewing ATA disk information

You can use the `disk shm_stats` command to view the SMART data from the ATA disks. In addition, you can use the `environment shelf_log` command to display the shelf logfiles, which are automatically attached to AutoSupport messages.

Command for managing the storage subsystem

You can use the `storage` command to do the following:

- Manage disks, SCSI and Fibre Channel host adapters, and all other components of the storage subsystem connected to your filer
- Enable or disable a host adapter
- View multiple host adapter paths to disks
- View information about hubs, disks, tape devices, and medium changer devices connected to your filer

The storage command syntax and subcommands

The `storage` command syntax is as follows:

```
storage <sub_command>
```

The subcommands are

```
alias [alias {electrical_name | wwn }]
disable adapter name
enable adapter name
help sub_command
show adapter [-a] [name ]
show hub [-a] [name ]
show disk [-a | -p]
show mc [name]]
show port [name]
show switch [name]
show tape [name]
show tape supported [-v]
stats tape name
stats tape zero name
unalias {alias | -a | -m | -t }
```

For detailed information about the `storage` command and all options available with this command, see the na_storage(1) man page on the filer. The options `alias` and `unalias` are discussed in detail in the *Data Protection Guide Tape Backup and Recovery Guide*.

Detailed information

For detailed information on how to perform specific tasks by using the `storage` command, see the following topics:

- "Changing the state of an adapter" on page 94
- "Viewing storage subsystem information" on page 96

Changing the state of an adapter

About the state of a adapter

An adapter can be enabled or disabled. You can change the state of an adapter by using the `storage` command.

When to change the state of an adapter

Disable: You might want to disable an adapter if

- You are replacing any of the hardware components connected to the adapter, such as cables and Gigabit Interface Converters (GBICs)
- You are replacing a malfunctioning LRC module or bad cables

You can disable an adapter only if all disks connected to it can be reached through another adapter. Consequently, SCSI adapters and adapters connected to single-attached devices cannot be disabled.

If you try to disable an adapter that is connected to disks with no redundant access paths, you will get the following error message:

```
"Some device(s) on host adapter n can only be
 accessed through this adapter; unable to disable adapter"
```

After an adapter connected to dual-connected disks has been disabled, the other adapter is not considered redundant; thus, the other adapter cannot be disabled.

Enable: You might want to enable a disabled adapter after you have performed maintenance.

Enabling or disabling an adapter

To enable or disable an adapter, complete the following steps.

Step	Action
1	Enter the following command to identify the name of the adapter whose state you want to change: **`storage show adapter`** **Result:** The field that is labeled "Slot" lists the adapter name.

Step	Action	
2	**If you want to...**	**Then...**
	Enable the adapter	Enter the following command: **`storage enable adapter`** ***`name`*** *name* is the adapter name.
	Disable the adapter	Enter the following command: **`storage disable adapter`** ***`name`*** *name* is the adapter name.

Viewing storage subsystem information

Command for viewing information

You can use the `storage show` command and its options to view information about the storage subsystem components connected to your filer, including the adapters, hubs, disks, tape devices, and medium changer devices.

For detailed information about the options available with the `storage show` command, see the na_storage(1) man page on the filer.

Viewing information about all components

To view information about storage subsystem components, complete the following step.

Step	Action
1	Enter the following command: **`storage show`**

Example: The following example shows information about the adapters and disks connected to the filer tpubs-cf1:

```
tpubs-cf1> storage show
Slot:            7
Description:     Fibre Channel Host Adapter 7 (QLogic 2100 rev. 3)
Firmware Rev:    1.19.14
PCI Bus Width:   32-bit
PCI Clock Speed: 33 MHz
FC Node Name:    2:000:00e08b:006a15
Cacheline Size:  128
FC Packet Size:  512
Link Data Rate:  2 Gbit
SRAM Parity:     No
External GBIC:   No
State:           Enabled
In Use:          Yes
Redundant:       No
```

```
DISK  SHELF  BAY  SERIAL       VENDOR   MODEL      REV
----- -----  ---  ------       ------   ---------  ----
7.6     0     6   LA774453     SEAGATE  ST19171FC  FB59
7.5     0     5   LA694863     SEAGATE  ST19171FC  FB59
7.4     0     4   LA781085     SEAGATE  ST19171FC  FB59
7.3     0     3   LA773189     SEAGATE  ST19171FC  FB59
7.14    1     6   LA869459     SEAGATE  ST19171FC  FB59
7.13    1     5   LA781479     SEAGATE  ST19171FC  FB59
7.12    1     4   LA772259     SEAGATE  ST19171FC  FB59
7.11    1     3   LA783073     SEAGATE  ST19171FC  FB59
7.10    1     2   LA700702     SEAGATE  ST19171FC  FB59
7.9     1     1   LA786084     SEAGATE  ST19171FC  FB59
7.8     1     0   LA761801     SEAGATE  ST19171FC  FB59
7.2     0     2   LA708093     SEAGATE  ST19171FC  FB59
7.1     0     1   LA773443     SEAGATE  ST19171FC  FB59
7.0     0     0   LA780611     SEAGATE  ST19171FC  FB59
```

Viewing information about host adapters

To view information about host adapters, complete the following step.

Step	Action	
1	**If you want to view...**	**Then...**
	Information about all the host adapters	Enter the following command: **`storage show adapter`**
	Information about a specific host adapter	Enter the following command: **`storage show adapter`** ***`name`*** *name* is the adapter name.

Example 1: The following example shows information about all the adapters installed in the filer tpubs-cf2:

```
tpubs-cf2> storage show adapter
Slot:            7a
Description:     Fibre Channel Host Adapter 7a (QLogic 2100 rev. 3)
Firmware Rev:    1.19.14
PCI Bus Width:   32-bit
PCI Clock Speed: 33 MHz
FC Node Name:    2:000:00e08b:00fb15
Cacheline Size:  128
FC Packet Size:  512
Link Data Rate:  2 Gbit
```

```
SRAM Parity:     No
External GBIC:   No
State:           Enabled
In Use:          Yes
Redundant:       No
Slot:            7b
Description:     Fibre Channel Host Adapter 7b (QLogic 2100 rev. 3)
Firmware Rev:    1.19.14
PCI Bus Width:   32-bit
PCI Clock Speed: 33 MHz
FC Node Name:    2:000:00e08b:006b15
Cacheline Size:  128
FC Packet Size:  512
Link Data Rate:  2 Gbit
SRAM Parity:     No
External GBIC:   No
State:           Enabled
In Use:          Yes
Redundant:       No
```

Example 2: The following example shows information about adapter 7b in the filer tpubs-cf2:

```
tpubs-cf2> storage show adapter 7b
Slot:            7b
Description:     Fibre Channel Host Adapter 7b (QLogic 2100 rev. 3)
Firmware Rev:    1.19.14
PCI Bus Width:   32-bit
PCI Clock Speed: 33 MHz
FC Node Name:    2:000:00e08b:006b15
Cacheline Size:  128
FC Packet Size:  512
Link Data Rate:  2 Gbit
SRAM Parity:     No
External GBIC:   No
State:           Enabled
In Use:          Yes
Redundant:       No
```

Viewing information about hubs

To view information about hubs, complete the following step.

Step	Action	
1	**If you want to view...**	**Then...**
	Information about all hubs	Enter the following command: **`storage show hub`**
	Information about a specific hub	Enter the following command: **`storage show hub`** ***`name`*** *name* is the hub name.

Example: The following example shows information about hub 8a.shelf1:

```
storage show hub 8a.shelf1
Hub name:    8a.shelf1
Channel:     8a
Loop:        A
Shelf id:    1
Shelf UID:   50:05:0c:c0:02:00:12:3d
Term switch: OFF
Shelf state: ONLINE
ESH state:   OK
                 Loop Invalid  Invalid Clock  Insert Stall Util
Disk  Disk   Port     up       CRC     Word   Delta  Count Count %
Id    Bay    State    Count    Count   Count
-----------------------------------------------------------------
[IN ]        OK   3    0        0      128    1      0     0
[ 16]  0     OK   4    0        0      128    0      0     0
[ 17]  1     OK   4    0        0      128    0      0     0
[ 18]  2     OK   4    0        0      128    0      0     0
[ 19]  3     OK   4    0        0      128    0      0     0
[ 20]  4     OK   4    0        0      128    0      0     0
[ 21]  5     OK   4    0        0      128    0      0     0
[ 22]  6     OK   4    0        0      128    0      0     0
[ 23]  7     OK   4    0        0      128    0      0     0
[ 24]  8     OK   4    0        0      128    0      0     0
[ 25]  9     OK   4    0        0      128    0      0     0
[ 26]  10    OK   4    0        0      128    0      0     0
[ 27]  11    OK   4    0        0      128    0      0     0
[ 28]  12    OK   4    0        0      128    0      0     0
[ 29]  13    OK   4    0        0      128    0      0     0
[OUT]        OK   4    0        0      128    0      0     0
```

```
Hub name:    8b.shelf1
Channel:     8b
Loop:        B
Shelf id:    1
Shelf UID:   50:05:0c:c0:02:00:12:3d
Term switch: OFF
Shelf state: ONLINE
ESH state:   OK
                 Loop  Invalid  Invalid Clock  Insert Stall Util
Disk  Disk   Port     up       CRC     Word   Delta  Count Count %
Id    Bay    State    Count    Count   Count
-------------------------------------------------------------------
[IN ]        OK   3    0        0       128    1      0     0
[ 16]  0     OK   4    0        0       128    0      0     0
[ 17]  1     OK   4    0        0       128    0      0     0
[ 18]  2     OK   4    0        0       128    0      0     0
[ 19]  3     OK   4    0        0       128    0      0     0
[ 20]  4     OK   4    0        0       128    0      0     0
[ 21]  5     OK   4    0        0       128    0      0     0
[ 22]  6     OK   4    0        0       128    0      0     0
[ 23]  7     OK   4    0        0       128    0      0     0
[ 24]  8     OK   4    0        0       128    0      0     0
[ 25]  9     OK   4    0        0       128    0      0     0
[ 26]  10    OK   4    0        0       128    0      0     0
[ 27]  11    OK   4    0        0       128    0      0     0
[ 28]  12    OK   4    0        0       128    0      0     0
[ 29]  13    OK   4    0        0       128    0      0     0
[OUT]        OK   4    0        0       128    0      0     0
```

Note

Hub 8b.shelf1 is also listed by the `storage show hub 8a.shelf1` command in the example, because the two hubs are part of the same shelf and the disks in the shelf are dual-ported disks. Effectively, the command is showing the disks from two perspectives.

Viewing information about medium changers

To view information about medium changers attached to your filer, complete the following step.

Step	Action
1	Enter the following command: **`storage show mc [`*`name`*`]`** *name* is the name of the medium changer for which you want to view information. If no medium changer name is specified, information for all medium changers is displayed.

Viewing information about switch ports

To view information about ports for switches attached to the filer, complete the following step.

Step	Action
1	Enter the following command: **`storage show port [`*`name`*`]`** *name* is the name of the port for which you want to view information. If no port name is specified, information for all ports is displayed.

Viewing information about switches

To view information about switches attached to the filer, complete the following step.

Step	Action
1	Enter the following command: **`storage show switch [`*`name`*`]`** *name* is the name of the switch for which you want to view information. If no switch name is specified, information for all switches is displayed.

Viewing information about tape drives

To view information about tape drives attached to your filer, complete the following step.

Step	Action
1	Enter the following command: **`storage show tape [`*`tape`*`]`** *tape* is the name of the tape drive for which you want to view information. If no tape name is specified, information for all tape drives is displayed.

Viewing supported tape drives

To view information about tape drives that are supported by your filer, complete the following step.

Step	Action
1	Enter the following command: **`storage show tape supported [-v]`** `-v` displays all information about supported tape drives, including their density and compression settings. If no option is given, only the names of supported tape drives are displayed.

Viewing tape drive statistics

To view storage statistics for tape drives attached to the filer, complete the following step.

Step	Action
1	Enter the following command: **`storage stats tape `*`name`*** *name* is the name of the tape drive for which you want to view storage statistics.

Resetting tape drive statistics

To reset storage statistics for a tape drive attached to the filer, complete the following step.

Step	Action
1	Enter the following command: `storage stats tape zero` *`name`* *name* is the name of the tape drive.

RAID Protection of Data 4

About this chapter

This chapter describes how to manage RAID protection on NetApp appliance aggregates. Throughout this chapter, aggregates refers to those that contain both FlexVol and traditional volumes.

Data ONTAP uses RAID Level 4 or RAID-DP (double-parity) protection to ensure data integrity within a group of disks even if one or two of those disks fail.

Note

The RAID principles and management operations described in this chapter do not apply to gFiler gateways. Data ONTAP uses RAID-0 for gFilers since the LUNs that they use are RAID protected by the storage subsystem.

Topics in this chapter

This chapter discusses the following topics:

- "Understanding RAID groups" on page 106
- "Predictive disk failure and Rapid RAID Recovery" on page 114
- "Disk failure and RAID reconstruction with a hot spare disk" on page 115
- "Disk failure without a hot spare disk" on page 116
- "Replacing disks in a RAID group" on page 118
- "Setting RAID type and group size" on page 119
- "Changing the RAID type for an aggregate" on page 122
- "Changing the size of RAID groups" on page 127
- "Controlling the speed of RAID operations" on page 130
- "Automatic and manual disk scrubs" on page 135
- "Minimizing media error disruption of RAID reconstructions" on page 142
- "Viewing RAID status" on page 150

Understanding RAID groups

About RAID groups in Data ONTAP

A RAID group consists of one or more data disks, across which client data is striped and stored, plus one or two parity disks. The purpose of a RAID group is to provide parity protection from data loss across its included disks. RAID-4 uses one parity disk to ensure data recoverability if one disk fails within the RAID group. RAID-DP uses two parity disks to ensure data recoverability even if two disks within the RAID group fail.

RAID group disk types

Data ONTAP assigns and makes use of four different disk types to support data storage, parity protection, and disk replacement.

Disk	Description
Data disk	Holds data stored on behalf of clients within RAID groups (and any data generated about the state of the filer as a result of a malfunction).
Hot spare disk	Does not hold usable data, but is available to be added to a RAID group in an aggregate. Any functioning disk that is not assigned to an aggregate functions as a hot spare disk.
Parity disk	Stores data reconstruction information within RAID groups.
dParity disk	Stores double-parity information within RAID groups, if RAID-DP is enabled.

Types of RAID protection

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

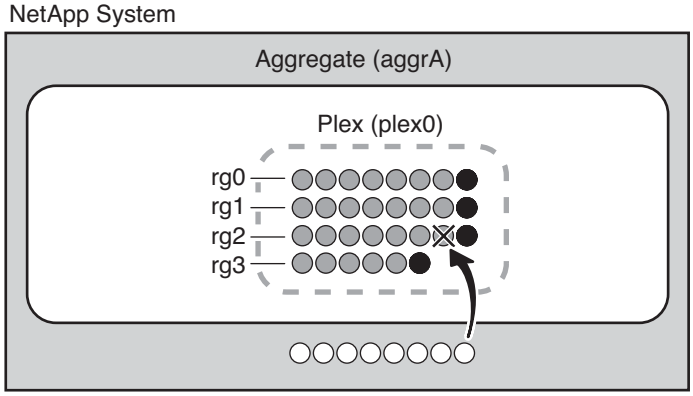
- If an aggregate is configured for RAID-4 protection, Data ONTAP reconstructs the data from a single failed disk within a RAID group and transfers that reconstructed data to a spare disk.
- If an aggregate is configured for RAID-DP protection, Data ONTAP reconstructs the data from one or two failed disks within a RAID group and transfers that reconstructed data to one or two spare disks as necessary.

RAID-4 protection: RAID-4 provides single-parity disk protection against single-disk failure within a RAID group. The minimum number of disks in a RAID 4 group is two: at least one data disk and one parity disk. If there is a single data or parity disk failure in a RAID-4 group, Data ONTAP replaces the failed disk in the RAID group with a spare disk and uses the parity data to reconstruct the failed disk's data on the replacement disk. If there are no spare disks available, Data ONTAP goes into a degraded mode and alerts you of this condition.

Caution

With RAID-4, if there is a second disk failure before data can be reconstructed from the data on the first failed disk, there will be data loss. To avoid data loss when two disks fail, you can select RAID-DP. This provides two parity disks to protect you from data loss when two disk failures occur in the same RAID group before the first failed disk can be reconstructed.

The following figure diagrams a traditional volume configured for RAID-4 protection.



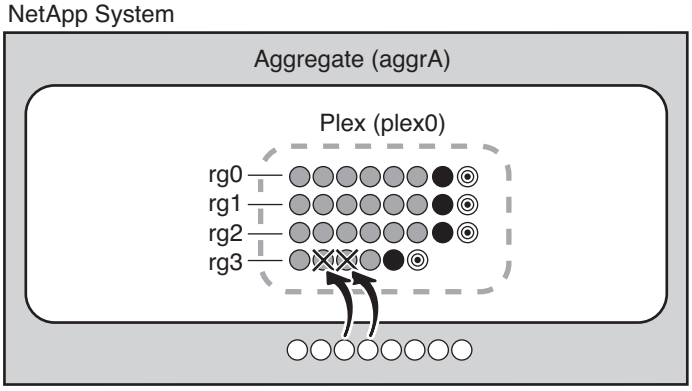
RAID-DP protection: RAID-DP provides double-parity disk protection when the following conditions occur:

- There are media errors on a block when Data ONTAP is attempting to reconstruct a failed disk.
- There is a single- or double-disk failure within a RAID group.

The minimum number of disks in a RAID-DP group is three: at least one data disk, one regular parity disk, and one double-parity (or dParity) disk.

If there is a data-disk or parity-disk failure in a RAID-DP group, Data ONTAP replaces the failed disk in the RAID group with a spare disk and uses the parity

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



How Data ONTAP organizes RAID groups automatically

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

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

  Note

  If an existing RAID-4 group is assigned an additional disk that is larger than the group's existing parity disk, then Data ONTAP reassigns the new disk as parity disk for that RAID group. If all disks are of equal size, any one of the disks can be selected for parity.

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

Note

If an existing RAID-DP group is assigned an additional disk that is larger than the group's existing dParity disk, then Data ONTAP reassigns the new disk as the regular parity disk for that RAID group and restricts its capacity to be no greater than that of the existing dParity disk. If all disks are of equal size, any one of the disks can be selected for the dParity disk.

Hot spare disks

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

Hot spare disk size recommendations: NetApp recommends keeping at least one hot spare disk for each disk size and disk type installed in your filer. This allows the filer to use a disk of the same size and type as a failed disk when reconstructing a failed disk. If a disk fails and a hot spare disk of the same size is not available, the filer uses a spare disk of the next available size up. See "Disk failure and RAID reconstruction with a hot spare disk" on page 115 for more information.

Note

If no spare disks exist in a filer, Data ONTAP can continue to function in degraded mode. Data ONTAP supports degraded mode in the case of single-disk failure for aggregates configured with RAID-4 protection and in the case of single- or double- disk failure in aggregates configured for RAID-DP protection. For details see "Disk failure without a hot spare disk" on page 116.

Maximum number of RAID groups

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

RAID-4, RAID-DP, and SyncMirror

RAID-4 and RAID-DP can be used in combination with the Data ONTAP SyncMirror feature, which also offers protection against data loss due to disk or other hardware component failure. SyncMirror protects against data loss by maintaining two copies of the data contained in the aggregate, one in each plex.

Any data loss due to disk failure in one plex is repaired by the undamaged data in the opposite plex. The advantages and disadvantages of using RAID-4 or RAID-DP, with and without the SyncMirror feature, are listed in the following tables.

Advantages and disadvantages of using RAID-4:

Factor affected by RAID type	**RAID-4**	**RAID-4 with SyncMirror**
What RAID and SyncMirror protect against	Single-disk failure within one or multiple RAID groups	Single-disk failure within one or multiple RAID groups in one plex and single-, double-, or greater-disk failure in the other plex. A double-disk failure in a RAID group results in a failed plex. If this occurs, a double-disk failure on any RAID group on the other plex fails the aggregate. See "Advantages of RAID-4 with SyncMirror" on page 111. Storage subsystem failures (HBA, cables, shelf) on the filer
Required disk resources per RAID group	*n* data disks + 1 parity disk	2 x (*n* data disks + 1 parity disk)
Performance cost	None	Low mirroring overhead; can improve performance
Additional cost and complexity	None	SyncMirror license and configuration Possible cluster license and configuration

Advantages and disadvantages of using RAID-DP:

Factor affected by RAID type	RAID-DP	RAID-DP with SyncMirror
What RAID and SyncMirror protect against	Single- or double-disk failure within one or multiple RAID groups	Single-disk failure and media errors on another disk. Single- or double-disk failure within one or multiple RAID groups in one plex and single-, double-, or greater disk failure in the other plex. SyncMirror and RAID-DP together cannot protect against more than two disk failures on both plexes. It can protect against more than two disk failures on one plex with up to two disk failures on the second plex. A triple disk failure in a RAID group results in a failed plex. If this occurs, a triple disk failure on any RAID group on the other plex will fail the aggregate. See "Advantages of RAID-DP with SyncMirror" on page 112. Storage subsystem failures (HBA, cables, shelf) on the filer
Required disk resources per RAID group	*n* data disks + 2 parity disks	2 x (*n* data disks + 2 parity disks)
Performance cost	Almost none	Low mirroring overhead; can improve performance
Additional cost and complexity	None	SyncMirror license and configuration Possible cluster license and configuration

Advantages of RAID-4 with SyncMirror: On SyncMirror-replicated aggregates using RAID-4, any combination of multiple disk failures within single RAID groups in one plex is restorable, as long as multiple disk failures are not concurrently occurring in the opposite plex of the mirrored aggregate.

Advantages of RAID-DP with SyncMirror: On SyncMirror-replicated aggregates using RAID-DP, any combination of multiple disk failures within single RAID groups in one plex is restorable, as long as concurrent failures of more than two disks are not occurring in the opposite plex of the mirrored aggregate.

For more SyncMirror information: For more information on the Data ONTAP SyncMirror feature, see the *Data Protection Online Backup and Recovery Guide*.

Larger versus smaller RAID groups

You can specify the number of disks in a RAID group and the RAID level of protection, or you can use the default for the specific appliance. Adding more data disks to a RAID group increases the striping of data across those disks, which typically improves I/O performance. However, with more disks, there is a greater risk that one of the disks might fail.

Configuring an optimum RAID group size for an aggregate requires a trade-off of factors. You must decide which factor—speed of recovery, assurance against data loss, or maximizing data storage space—is most important for the aggregate that you are configuring. For a list of default and maximum RAID group sizes, see "Maximum and default RAID group sizes" on page 127.

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

- More data drives available. An aggregate configured into a few large RAID groups requires fewer drives reserved for parity than that same aggregate configured into many small RAID groups.
- Small improvement in system performance. Write operations are generally faster with larger RAID groups than with smaller RAID groups.

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

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

For example, whether you have a RAID group with fourteen disks or two RAID groups with seven disks, you still have the same number of disks available for striping. However, with multiple smaller RAID groups, you minimize the risk of the performance impact during reconstruction and you minimize the risk of multiple disk failure within each RAID group.

Advantages of RAID-DP over RAID-4

With RAID-DP, you can use larger RAID groups because they offer more protection. A RAID-DP group is more reliable than a RAID-4 group that is half its size, even though a RAID-DP group has twice as many disks. Thus, the RAID-DP group provides better reliability with the same parity overhead.

Predictive disk failure and Rapid RAID Recovery

How Data ONTAP handles failing disks

Data ONTAP monitors disk performance so that when certain conditions occur, it can predict that a disk is likely to fail. For example, under some circumstances, if 100 or more media errors occur on a disk in a one-week period. When this occurs, Data ONTAP implements a process called Rapid RAID Recovery, and performs the following tasks:

1. Places the disk in question in pre-fail mode. This can occur at any time, regardless of what state the RAID group containing the disk is in.
2. Swaps in the spare replacement disk.
3. Copies the pre-failed disk's contents to a hot spare disk on the filer before an actual failure occurs.
4. Once the copy is complete, fails the disk that is in pre-fail mode.

Steps 2 through 4 can only occur when the RAID group is in a normal state.

By executing a copy, fail, and disk swap operation on a disk that is predicted to fail, Data ONTAP avoids three problems that a sudden disk failure and subsequent RAID reconstruction process involves:

- Rebuild time
- Performance degradation
- Potential data loss due to additional disk failure during reconstruction

If the disk that is in pre-fail mode fails on its own before copying to a hot spare disk is complete, Data ONTAP starts the normal RAID reconstruction process.

Disk failure and RAID reconstruction with a hot spare disk

About this section

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

Data ONTAP replaces failed disk with spare and reconstructs data

If a disk fails, Data ONTAP performs the following tasks:

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

Note

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

During reconstruction, file service might slow down.

Caution

After Data ONTAP is finished reconstructing data, NetApp recommends that you replace the failed disk or disks with new hot spare disks as soon as possible, so that hot spare disks are always available in the system. For information about replacing a disk, see Chapter 3, “Disk and Storage Subsystem Management,” on page 43.

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

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

Disk failure without a hot spare disk

About this section

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

Filer runs in degraded mode

When there is a single-disk failure in RAID-4 enabled aggregates or a double-disk failure in RAID-DP enabled aggregates, and there are no hot spares available, the filer continues to run without losing any data, but performance is somewhat degraded.

Caution

You should replace the failed disks as soon as possible, because additional disk failure might cause the filer to lose data in the file systems contained in the affected aggregate.

Filer logs warning messages in /etc/messages

The filer logs a warning message in the /etc/messages file on the root volume once per hour after a disk fails.

Filer shuts down automatically after 24 hours

To ensure that you notice the failure, the filer automatically shuts itself off in 24 hours, by default, or at the end of a period that you set with the `raid.timeout` option of the `options` command. You can restart the filer without fixing the disk, but it continues to shut itself off periodically until you repair the problem.

Filer sends messages about failures

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

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

Filer reconstructs data after disk is replaced

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

Replacing disks in a RAID group

Replacing data disks

If you need to replace a disk—for example a mismatched data disk in a RAID group—you use the `disk replace` command. This command uses Rapid RAID Recovery to copy data from the specified old disk in a RAID group to the specified spare disk in the appliance. At the end of the process, the spare disk replaces the old disk as the new data disk, and the old disk becomes a spare disk in the appliance.

Note

Data ONTAP does not allow mixing disk types in the same aggregate.

To replace a disk in a RAID group, complete the following step.

Step	Action
1	Enter the following command: `disk replace start [-f] old_disk new_spare` `-f` suppresses confirmation information being displayed. It also allows a less than optimum replacement disk to be used. For example, the replacement disk might not have a matching RPM, or it might not be in the right spare pool.

Stopping the disk replacement operation

To stop the disk replace operation, or to prevent the operation if copying did not begin, complete the following step.

Step	Action
1	Enter the following command: `disk replace stop old_disk`

Setting RAID type and group size

About RAID group type and size

Data ONTAP provides default values for the RAID group type and RAID group size parameters when you create aggregates and traditional volumes. You can use these defaults or you can specify different values.

Specifying the RAID type and size when creating aggregates or FlexVol volumes

To specify the type and size of an aggregate's or traditional volume's RAID groups, complete the following steps.

Step	Action
1	View the spare disks to know which ones are available to put in a new aggregate by entering the following command: `aggr status -s` **Result:** The device number, shelf number, and capacity of each spare disk on the filer is listed.

2	For an aggregate, specify RAID group type and RAID group size by entering the following command:
	`aggr create aggr [-m] [-t {raid4|raid_dp}] [-r raid_group_size] disk_list`
	aggr is the name of the aggregate you want to create.
	or
	For a traditional volume, specify RAID group type and RAID group size by entering the following command:
	`aggr create vol [-v] [-m] [-t {raid4|raid_dp}] [-r raid_group_size] disk_list`
	vol is the name of the traditional volume you want to create.

Step	Action
	`-m` specifies the optional creation of a SyncMirror-replicated volume if you want to supplement RAID protection with SyncMirror protection. A SyncMirror license is required for this feature. `-t {raid4|raid_dp}` specifies the type of RAID protection (RAID-4 or RAID-DP) that you want to provide. If no RAID type is specified, the default value `raid_dp` is applied to an aggregate or the default value `raid4` is applied to a traditional volume. RAID-DP is the default for both aggregates and traditional volumes on NearStore systems that support ATA disks. `-r` *raid-group-size* is the number of disks per RAID group that you want. If no RAID group size is specified, the default value for your appliance model is applied. For a listing of default and maximum RAID group sizes, see "Maximum and default RAID group sizes" on page 127. *disk_list* specifies the disks to include in the volume that you want to create. It can be expressed in the following formats: ◆ *ndisks*[*@disk-size*] *ndisks* is the number of disks to use. It must be at least 2. *disk-size* is the disk size to use, in gigabytes. You must have at least *ndisks* available disks of the size you specify. ◆ `-d` *disk_name1 disk_name2... disk_nameN* *disk_name1*, *disk_name2*, and *disk_nameN* are disk IDs of one or more available disks; use a space to separate multiple disks. **Example:** The following command creates the aggregate newaggr. Since RAID-DP is the default, it does not have to be specified. RAID group size is 16 disks. Since the aggregate consists of 32 disks, those disks will form two RAID groups, rg0 and rg1: `aggr create newaggr -r 16 32@72`

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

Changing the RAID type for an aggregate

Changing the RAID group type

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

Changing from RAID-4 to RAID-DP protection

Before you change an aggregate's RAID protection from RAID-4 to RAID-DP, you need to ensure that hot spare disks of sufficient number and size are available. During the conversion, Data ONTAP adds an additional disk to each existing RAID group from the filer's hot spare disk pool and assigns the new disk the dParity disk function for the RAID-DP group. In addition, the aggregate's `raidsize` option is changed to RAID-DP as the default on this appliance. The `raidsize` option also controls the size of new RAID groups that might be created in the aggregate.

Changing an aggregate's RAID type: To change an existing aggregate's RAID protection from RAID-4 to RAID-DP, complete the following steps.

Step	Action
1	Determine the number of RAID groups and the size of their parity disks in the aggregate in question by entering the following command. **`aggr status` *`aggr_name`* `-r`**
2	Enter the following command to make sure that a hot spare disk exists on the filer for each RAID group listed for the aggregate in question, and make sure that these hot spare disks match the size and checksum type of the existing parity disks in those RAID groups. **`aggr status -s`** If necessary, add hot spare disks of the appropriate number of appropriate number, size, and checksum type to the filer. See "Adding new disks" on page 63.

Step	Action
3	Enter the following command: **aggr options *aggr_name* raidtype raid_dp** *aggr_name* is the aggregate whose RAID type you are changing. **Example:** The following command changes the RAID type of the aggregate this aggr to RAID-DP: `aggr options thisaggr raidtype raid_dp` For backward compatibility, you can enter the following command: **vol options *vol_name* raidtype raid_dp**

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

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

  If hot spare disks available on the filer are of insufficient number or size to support the RAID type conversion, Data ONTAP issues a warning before executing the command to set the RAID type to RAID-DP (either `aggr options` *aggr_name* `raidtype raid_dp` or `vol options` *vol_name* `raidtype raid_dp`).

  If you continue the operation, RAID-DP protection is implemented on the aggregate in question, but some of its RAID groups for which no second parity disk was available remain degraded. In this case, the protection offered is no improvement over RAID-4, and no hot spare disks are available in case of disk failure since all were reassigned as dParity disk.
- The aggregate's `raidsize` option, which sets the size for any new RAID groups created in this aggregate, is automatically reset to one of the following RAID-DP defaults:
  - On all non-NearStore appliances, 16
  - On an R100 platform, 12
  - On an R150 platform, 12
  - On an R200 platform, 14

Note

After the `aggr options` *aggr_name* `raidtype raid_dp` operation is complete, you can manually change the `raidsize` option through the `aggr options` *aggr_name* `raidsize` command. See "Changing the maximum size of RAID groups" on page 128.

For backward compatibility, you can also use the following commands for traditional volumes:
`vol options` *vol_name* `raidtype raid_dp` operation
`vol options` *vol_name* `raidsize` command

Changing from RAID-DP to RAID-4 protection

Changing an aggregate's RAID type: While it is possible to change an aggregate from RAID-DP to RAID-4, there is a restriction, as described in the following note.

Note

You cannot change an aggregate from RAID-DP to RAID 4 if the aggregate contains a RAID group larger than the maximum allowed for RAID 4.

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

Step	Action
1	Enter the following command: **`aggr options aggr_name raidtype raid4`** *aggr_name* is the aggregate whose RAID type you are changing. **Example:** The following command changes the RAID type of the aggregate thataggr to RAID-4: `aggr options thataggr raidtype raid4`

Associated RAID group size changes: The RAID group size determines the size of any new RAID groups created in an aggregate. When you change the RAID protection of an existing aggregate from RAID-DP to RAID-4, Data ONTAP automatically carries out the following associated RAID group size changes:

- In each of the aggregate's existing RAID groups, the RAID-DP second parity disk (dParity) is removed and placed in the hot spare disk pool, thus reducing each RAID group's size by one parity disk.
- For NearStore appliances, Data ONTAP changes the aggregate's `raidsize` option to the RAID-4 default sizes, as indicated on the following platforms:
  - R100 (8)
  - R150 (6)
  - R200 (7)
- For non-NearStore appliances, Data ONTAP changes the setting for the aggregate's `raidsize` option to the size of the largest RAID group in the aggregate. However, there are two exceptions:
  - If the aggregate's largest RAID group is larger than the maximum RAID-4 group size on non-NearStore platforms (14), then the aggregate's `raidsize` option is set to 14.
  - If the aggregate's largest RAID group is smaller than the default RAID-4 group size on non-NearStore platforms (8), then the aggregate's `raidsize` option is set to 8.

After the `aggr options` *aggr_name* `raidtype raid_dp` operation is complete, you can manually change the `raidsize` option through the `aggr options` *aggr_name* `raidsize` command. See "Changing the maximum size of RAID groups" on page 128.

For backward compatibility, you can also use the following commands for traditional volumes:

`vol options` *vol_name* `raidtype raid_dp`
`vol options` *vol_name* `raidsize`

Verifying the RAID type

To verify the RAID type of an aggregate, complete the following step.

Step	Action
1	Enter the following command: `aggr status` *`aggr_name`* or `aggr options` *`aggr_name`* For backward compatibility, you can also enter the following command: `vol options` *`vol_name`*

Changing the size of RAID groups

Maximum and default RAID group sizes

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

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

Maximum and default RAID-DP group sizes and defaults: The following table lists the minimum, maximum, and default RAID-DP group sizes supported on NetApp platforms.

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

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

Platform	Minimum group size	Maximum group size	Default group size
R200	2	7	7
R150	2	6	6
R100	2	8	8
FAS250	2	14	7
All other NetApp platforms	2	14	8

Note

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

Changing the maximum size of RAID groups

The `aggr option raidsize` option specifies the maximum RAID group size that can be reached by adding disks to an aggregate. For backward compatibility, you can also use the `vol option raidsize` option when you change the `raidsize` option of a traditional volume's containing aggregate.

- You can increase the `raidsize` option to allow more disks to be added to the most recently created RAID group.
- The new `raidsize` setting also applies to subsequently created RAID groups in an aggregate. Either increasing or decreasing `raidsize` settings will apply to future RAID groups.
- You cannot decrease the size of already created RAID groups.
- Existing RAID groups remain the same size they were before the `raidsize` setting was changed.

Changing the raidsize setting: To change the `raidsize` setting for an existing aggregate, complete the following step.

Step	Action

1	Enter the following command:
	`aggr options aggr_name raidsize size`
	aggr_name is the aggregate whose `raidsize` setting you are changing.
	size is the number of disks you want in the most recently created and all future RAID groups in this aggregate.
	Example: The following command changes the `raidsize` setting of the aggregate yeraggr to 16 disks:
	`aggr options yeraggr raidsize 16`
	For backward compatibility, you can also enter the following command for traditional volumes:
	`vol options vol_name raidsize size`
	Example: The following command changes the `raidsize` setting of the traditional volume yervol to 16 disks:
	`vol options yervol raidsize 16`

For information about adding disks to existing RAID groups, see "Adding disks to aggregates" on page 167.

Verifying the raidsize setting

To verify the size of `raidsize` setting in an aggregate, enter the `aggr options` *`aggr_name`* command.

For backward compatibility, you can also enter the `vol options` *`vol_name`* command for traditional volumes.

Changing the size of existing RAID groups

If you increased the `raidsize` setting for an aggregate or a traditional volume, you can also use the `-g` *raidgroup* option in the `aggr add` command or in the `vol add` command to add disks to an existing RAID group. For information about adding disks to existing RAID groups, see "Adding disks to a specific RAID group in an aggregate" on page 170.

Controlling the speed of RAID operations

RAID operations you can control

You can control the speed of the following RAID operations with RAID options:

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

Effects of varying the speed on filer performance

The speed that you select for each of these operations might affect the overall performance of the filer. However, if the operation is already running at the maximum speed possible and it is fully utilizing one of the three system resources (the CPU, disks, or the FC loop on FC-based systems), changing the speed of the operation has no effect on the performance of the operation or the filer.

If the operation is not yet running, you can set a speed that minimally slows filer network operations or a speed that severely slows filer network operations. For each operation, NetApp recommends the following guidelines:

- If you want to reduce the performance impact that the operation has on client access to the filer, change the specific RAID option from `medium` (the default) to `low`. This also causes the operation to slow down.
- If you want to speed up the operation, change the RAID option from `medium` to `high`. This might decrease the performance of the filer in response to client access.

Detailed information

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

- "Controlling the speed of RAID data reconstruction" on page 131
- "Controlling the speed of disk scrubbing" on page 132
- "Controlling the speed of plex resynchronization" on page 133
- "Controlling the speed of mirror verification" on page 134

Controlling the speed of RAID data reconstruction

About RAID data reconstruction

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

Changing RAID data reconstruction speed

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

Step	Action
1	Enter the following command: **`options raid.reconstruct.perf_impact`** ***`impact`*** *impact* can be `high`, `medium`, or `low`. High means that the filer uses most of the system resources—CPU time, disks, and FC loop bandwidth (on FC-based systems)—available for RAID data reconstruction; this setting can heavily affect filer performance. Low means that the filer uses very little of the system resources; this setting lightly affects filer performance. The default speed is medium. **Note** The setting for this option also controls the speed of Rapid RAID recovery.

RAID operations affecting RAID data reconstruction speed

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

Controlling the speed of disk scrubbing

About disk scrubbing

Disk scrubbing means periodically checking the disk blocks of all disks on the filer for media errors and parity consistency.

Although disk scrubbing slows the filer somewhat, network clients might not notice the change in filer performance because disk scrubbing starts automatically at 1:00 a.m. on Sunday by default, when most systems are lightly loaded, and stops after six hours. You can change the start time with the `scrub sched` option, and you can change the duration time with the `scrub duration` option.

Changing disk scrub speed

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

Step	Action
1	Enter the following command: **`options raid.scrub.perf_impact`** ***impact*** *impact* can be `high`, `medium`, or `low`. High means that the filer uses most of the system resources—CPU time, disks, and FC loop bandwidth (on FC-based systems)—available for disk scrubbing; this setting can heavily affect filer performance. Low means that the filer uses very little of the system resources; this setting lightly affects filer performance. The default speed is low.

RAID operations affecting disk scrub speed

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

Controlling the speed of plex resynchronization

What plex resynchronization is

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

When plex resynchronization occurs

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

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

Changing plex resynchronization speed

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

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

RAID operations affecting plex resynchronization speed

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

Controlling the speed of mirror verification

What mirror verification is

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

Changing mirror verification speed

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

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

RAID operations affecting mirror verification speed

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

Automatic and manual disk scrubs

About disk scrubbing

Disk scrubbing means checking the disk blocks of all disks on the filer for media errors and parity consistency. If Data ONTAP finds media errors or inconsistencies, it fixes them by reconstructing the data from other disks and rewriting the data. Disk scrubbing reduces the chance of potential data loss as a result of media errors during reconstruction.

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

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

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

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

This section includes the following topics:

- "Scheduling an automatic disk scrub" on page 136
- "Manually running a disk scrub" on page 139

Scheduling an automatic disk scrub

About disk scrub scheduling

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

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

Rescheduling disk scrubbing

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

Step	Action
1	Enter the following command: **options raid.scrub.schedule *duration*{h|m}@*weekday*@*start_time* [,*duration*{h|m}@*weekday*@*start_time*] ...** *duration* {h|m} is the amount of time, in hours (h) or minutes (m) that you want to allot for this operation. **Note** If no duration is specified for a given scrub, the value specified in the raid.scrub.duration option is used. For details, see "Setting the duration of automatic disk scrubbing" on page 138. *weekday* is the day of the week (sun, mon, tue, wed, thu, fri, sat) when you want the operation to start. *start_time* is the hour of the day you want the scrub to start. Acceptable values are 0-23, where 0 is midnight and 23 is 11 p.m. **Example:** The following command schedules two weekly RAID scrubs. The first scrub is for four hours every Tuesday starting at 2 a.m. The second scrub is for eight hours every Saturday starting at 10 p.m. **options raid.scrub.schedule 240m@tue@2,8h@sat@22**
2	Verify your modification with the following command: **options raid.scrub.schedule** The duration, weekday, and start times for all your scheduled disk scrubs appear. **Note** If you want to restore the default automatic scrub schedule of Sunday at 1:00 a.m., reenter the command with an empty value string as follows: **options raid.scrub.schedule " "**.

Toggling automatic disk scrubbing

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

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

Setting the duration of automatic disk scrubbing

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

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

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

Note

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

Changing disk scrub speed

To change the speed of disk scrubbing, see “Controlling the speed of disk scrubbing” on page 132.

Manually running a disk scrub

About disk scrubbing and checking RAID group parity

You can manually run disk scrubbing to check RAID group parity on RAID groups at the RAID group level, plex level, or aggregate level. The parity checking function of the disk scrub compares the data disks in a RAID group to the parity disk in a RAID group. If during the parity check Data ONTAP determines that parity is incorrect, Data ONTAP corrects the parity disk contents.

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

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

Running manual disk scrubs on all aggregates

To run manual disk scrubs on all aggregates, complete the following step.

Step	Action
1	Enter the following command: **`aggr scrub start`**

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

Disk scrubs on specific RAID groups

To run a manual disk scrub on the RAID groups of a specific aggregate, plex, or RAID group, complete the following step.

Step	Action
1	Enter one of the following commands: **aggr scrub start *name*** *name* is the name of the aggregate, plex, or RAID group.

Examples:

In this example, the command starts the manual disk scrub on all the RAID groups in the aggr2 aggregate:

```
aggr scrub start aggr2
```

In this example, the command starts a manual disk scrub on all the RAID groups of plex1 of the aggr2 aggregate:

```
aggr scrub start aggr2/plex1
```

In this example, the command starts a manual disk scrub on RAID group 0 of plex1 of the aggr2 aggregate:

```
aggr scrub start aggr2/plex1/rg0
```

Stopping manual disk scrubbing

You might need to stop Data ONTAP from running a manual disk scrub. To stop a manual disk scrub, complete the following step. If you stop a disk scrub, you can not resume it at the same location. You must start the scrub from the beginning.

Step	Action
1	Enter the following command: **aggr scrub stop *aggr_name*** If *aggr_name* is not specified, Data ONTAP stops all manual disk scrubbing.

Suspending a manual disk scrub

Rather than stopping Data ONTAP from checking and correcting parity, you can suspend checking for any period of time and resume it later, at the same offset when you suspended the scrub.

To suspend manual disk scrubbing, complete the following step.

Step	Action
1	Enter the following commands: **`aggr scrub suspend`** ***`aggr_name`*** If *aggr_name* is not specified, Data ONTAP suspends all manual disk scrubbing.

Resuming a suspended disk scrub

To resume manual disk scrubbing, complete the following step.

Step	Action
1	Enter the following command: **`aggr scrub resume`** ***`aggr_name`*** If *aggr_name* is not specified, Data ONTAP resumes all suspended manual disk scrubbing.

Viewing disk scrub status

The disk scrub status tells you what percentage of the disk scrubbing has been completed. Disk scrub status also displays whether disk scrubbing of a volume, plex, or RAID group is suspended.

To view the status of a disk scrub, complete the following step.

Step	Action
1	Enter one of the following commands: **`aggr scrub status`** ***`aggr_name`*** If *aggr_name* is not specified, Data ONTAP shows the disk scrub status of all RAID groups.

Minimizing media error disruption of RAID reconstructions

About media error disruption prevention

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

- Improved handling of media errors by a WAFL repair mechanism. See "Handling of media errors during RAID reconstruction" on page 143.
- Default continuous media error scrubbing on filer disks. See "Continuous media scrub" on page 144.
- Continuous monitoring of disk media errors and automatic failing and replacement of disks that exceed system-defined media error thresholds. See "Disk media error failure thresholds" on page 149.

Handling of media errors during RAID reconstruction

About media error handling during RAID reconstruction

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

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

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

Purpose of the raid.reconstruction.wafliron.enable option

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

Recommendation: NetApp recommends that you leave the `raid.reconstruction.wafliron.enable` option at its default setting of On. Doing so might increase data availability.

Enabling and disabling the automatic wafliron process

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

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

Continuous media scrub

About continuous media scrubbing

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

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

Note

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

Not a substitute for a scheduled disk scrub: Because the continuous process described in this section scrubs only media errors, NetApp strongly recommends that you continue to run the filer's scheduled complete disk scrub operation, which is described in "Automatic and manual disk scrubs" on page 135. The complete disk scrub carries out parity and checksum checking and repair operations, in addition to media checking and repair operations, on a scheduled rather than a continuous basis.

Adjusting maximum time for a media scrub cycle

You can decrease the CPU resources consumed by a continuous media scrub under a heavy client workload by increasing the maximum time allowed for a media scrub cycle to complete from the default 72 hours.

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

If you want the media scrub operation to consume even fewer CPU resources under heavy filer client workload, you can increase the maximum number of hours for the media scrub cycle. This uses fewer CPU resources for the media scrub in times of heavy filer usage.

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

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

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

Disabling continuous media scrubbing

NetApp recommends that you keep continuous media error scrubbing enabled, particularly for R100 and R200 series systems, but you might decide to disable your continuous media scrub if your filer is carrying out operations with heavy performance impact and if you have alternative measures (such as aggregate SyncMirror replication or RAID-DP configuration) in place that prevent data loss due to filer disruption or double-disk failure.

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

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

Note

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

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

Checking media scrub activity

You can confirm media scrub activity on your filer by completing the following step.

Step	Action
1	Entering one of the following commands: **`aggr media_scrub status [/aggr[/plex][/raidgroup]] [-v]`** **`aggr media_scrub status [-s spare_disk_name] [-v]`** /*aggr*[/*plex*] [/*raidgroup*] is the optional pathname to the aggregate, plex, or RAID group on which you want to confirm media scrubbing activity. `-s` *disk_name* specifies the optional name of a specific spare disk on which you want to confirm media scrubbing activity. `-v` is specifies the verbose version of the media scrubbing activity status. The verbose status information includes the percentage of the current scrub that is complete, the start time of the current scrub, and the completion time of the last scrub. **Note** If you enter `aggr media_scrub` status without specifying a pathname or a disk name, the status of the current media scrubs on all RAID groups and spare disks is displayed.

Example 1. Checking of filer-wide media scrubbing: The following command displays media scrub status information for all the aggregates and spare disks on the filer.

```
aggr media_scrub status
aggr media_scrub /aggr0/plex0/rg0 is 0% complete
aggr media_scrub /aggr2/plex0/rg0 is 2% complete
aggr media_scrub /aggr2/plex0/rg1 is 2% complete
aggr media_scrub /aggr3/plex0/rg0 is 30% complete
aggr media_scrub 9a.8 is 31% complete
aggr media_scrub 9a.9 is 31% complete
aggr media_scrub 9a.13 is 31% complete
aggr media_scrub 9a.2 is 31% complete
aggr media_scrub 9a.12 is 31% complete
```

Example 2. Verbose checking of filer-wide media scrubbing: The following command displays verbose media scrub status information for all the aggregates on the filer.

```
aggr media_scrub status -v
aggr media_scrub: status of /aggr0/plex0/rg0 :
  Current instance of media_scrub is 0% complete.
  Media scrub started at Thu Mar  4 21:26:00 GMT 2004
  Last full media_scrub completed: Thu Mar  4 21:20:12 GMT 2004

aggr media_scrub: status of 9a.8 :
  Current instance of media_scrub is 31% complete.
  Media scrub started at Thu Feb 26 23:14:00 GMT 2004
  Last full media_scrub completed: Wed Mar  3 23:22:33 GMT 2004

aggr media_scrub: status of 9a.9 :
  Current instance of media_scrub is 31% complete.
  Media scrub started at Thu Feb 26 23:14:00 GMT 2004
  Last full media_scrub completed: Wed Mar  3 23:22:33 GMT 2004

aggr media_scrub: status of 9a.13 :
  Current instance of media_scrub is 31% complete.
  Media scrub started at Thu Feb 26 23:14:00 GMT 2004
  Last full media_scrub completed: Wed Mar  3 23:22:37 GM
```

Example 3. Checking for media scrubbing on a specific aggregate:

The following command displays media scrub status information for the aggregate *aggr2*.

```
aggr media_scrub status /aggr2
  aggr media_scrub /aggr2/plex0/rg0 is 4% complete
  aggr media_scrub /aggr2/plex0/rg1 is 10% complete
```

Example 4. Checking for media scrubbing on a specific spare disk:

The following commands display media scrub status information for the spare disk 9b.12.

```
aggr media_scrub status -s 9b.12
aggr media_scrub 9b.12 is 31% complete
```

```
aggr media_scrub status -s 9b.12 -v
aggr media_scrub: status of 9b.12 :
  Current instance of media_scrub is 31% complete.
  Media scrub started at Thu Feb 26 23:14:00 GMT 2004
  Last full media_scrub completed: Wed Mar  3 23:23:33 GMT 2004
```

Enabling continuous media scrubbing on spare disks

The continuous media scrub that is enabled by the system-wide default option setting, `options raid.media_scrub.enable on`, carries out its scrubbing operations only on appliance disks that have been assigned to an aggregate. An additional default option setting, `options raid.media_scrub.spares.enable on`, carries out the media scrubbing operation on an appliance's spare disks. For media scrubbing of spare disks to take place, both the system-wide option, `options raid.media_scrub.enable`, and the spare disk-specific option, `options raid.media_scrub.spares.enable`, must be set to On.

Disk media error failure thresholds

About media error thresholds

To prevent a filer panic or data loss that might occur if too many media errors are encountered during single-disk failure reconstruction, Data ONTAP tracks media errors on each active filer disk and sends a disk failure request to the RAID system if system-defined media error thresholds are crossed on that disk.

Disk media error thresholds that trigger an immediate disk failure request include

- More than twenty-five media errors (that are not related to disk scrub activity) occurring on a disk within a ten-minute period
- Three or more media errors occurring on the same sector of a disk

If the aggregate is not already running in degraded mode due to single-disk failure reconstruction when the disk failure request is received, Data ONTAP fails the disk in question, swaps in a hot spare disk, and begins RAID reconstruction to replace the failed disk.

In addition, if one hundred or more media errors occur on a disk in a one-week period, Data ONTAP pre-fails the disk and causes Rapid RAID Recovery to start. For more information, see “Predictive disk failure and Rapid RAID Recovery” on page 114.

Failing disks at the thresholds listed in this section greatly decreases the likelihood of a filer panic or double-disk failure during a single-disk failure reconstruction.

Viewing RAID status

About RAID status

You use the `aggr status` command to check the current RAID status and configuration for your aggregates.

To view RAID status for your aggregates, complete the following step.

Step	Action
1	Enter the following command: **`aggr status [`*`aggr_name`*`] -r`** *aggregate_name* is the name of the aggregate whose RAID status you want to view. **Note** If you omit the name of the aggregate (or the traditional volume), Data ONTAP displays the RAID status of all the aggregates on the system.

Possible RAID status displayed

The `aggr status -r` or `volume status -r` command displays the following possible status conditions that pertain to RAID:

- Degraded—The aggregate contains at least one degraded RAID group that is not being reconstructed after single- disk failure.
- Double degraded—The aggregate contains at least one RAID group with double-disk failure that is not being reconstructed (this state is possible if RAID-DP protection is enabled for the affected aggregate).
- Double reconstruction xx% complete—At least one RAID group in the aggregate is being reconstructed after experiencing a double-disk failure (this state is possible if RAID-DP protection is enabled for the affected aggregate).
- Mirrored—The aggregate is mirrored, and all of its RAID groups are functional.
- Mirror degraded—The aggregate is mirrored, and one of its plexes is offline or resynchronizing.
- Normal—The aggregate is unmirrored, and all of its RAID groups are functional.

- Partial—At least one disk was found for the aggregate, but two or more disks are missing.
- Reconstruction xx% complete—At least one RAID group in the aggregate is being reconstructed after experiencing a single- disk failure.
- Resyncing—The aggregate contains two plexes; one plex is resynchronizing with the aggregate.

Aggregate Management 5

About this chapter

This chapter describes how to use aggregates to manage filer storage resources.

Topics in this chapter

This chapter discusses the following topics:

- "Understanding aggregates" on page 154
- "Creating aggregates" on page 157
- "Changing the state of an aggregate" on page 162
- "Adding disks to aggregates" on page 167
- "Destroying aggregates" on page 172
- "Physically moving aggregates" on page 174

Understanding aggregates

Aggregate management

To support the differing security, backup, performance, and data sharing needs of your users, you can group the physical data storage resources on your appliance into one or more aggregates.

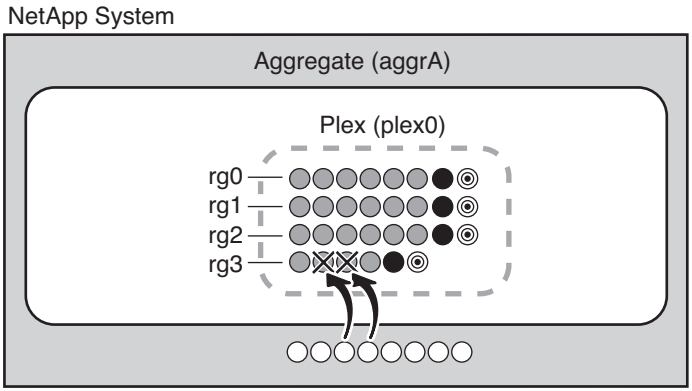
Each aggregate possesses its own RAID configuration, plex structure, and set of assigned disks. Within each aggregate you can create one or more FlexVol volumes—the logical file systems that share the physical storage resources, RAID configuration, and plex structure of that common containing aggregate.

For example, you can create a large aggregate with large numbers of disks in large RAID groups to support multiple FlexVol volumes, maximize your data resources, provide the best performance, and accommodate SnapVault backup.

You can also create a smaller aggregate to support FlexVol volumes that require special functions like SnapLock non-erasable data storage.

An unmirrored aggregate: In the following diagram, the unmirrored aggregate, arbitrarily named *aggrA* by the user, consists of one plex, which is made up of three double-parity RAID groups, automatically named *rg0*, *rg1*, and *rg2* by Data ONTAP.

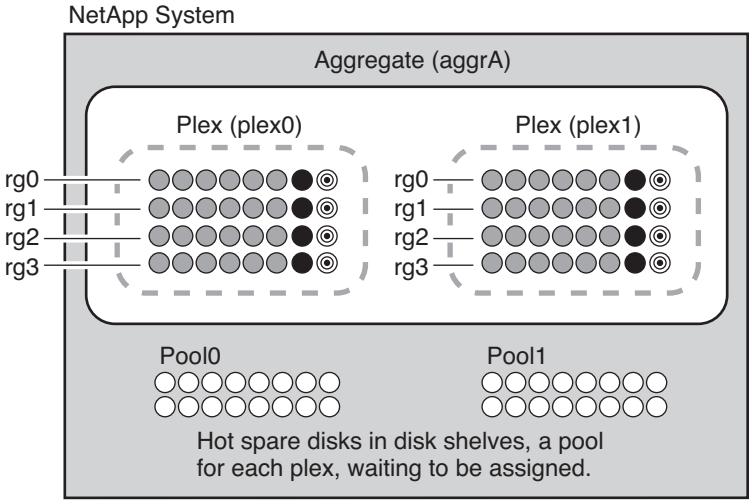
Notice that RAID-DP requires that both a parity disk and a double parity disk be in each RAID group. In addition to the disks that have been assigned to RAID groups, there are eight hot spare disks in the pool. In this diagram, two of the disks are needed to replace two failed disks, so only six will remain in the pool.



A mirrored aggregate: Consists of two plexes, which provides an even higher level of data redundancy via RAID-level mirroring. For an aggregate to be enabled for mirroring, the appliance must have a SyncMirror license for syncmirror_local or cluster_remote installed and enabled, and the appliance's disk configuration must support RAID-level mirroring.

When SyncMirror is enabled, all the disks are divided into two disk pools, and a copy of the plex is created. The plexes are physically separated (each plex has its own RAID groups and its own disk pool), and the plexes are updated simultaneously. This provides added protection against data loss if there is a double-disk failure or a loss of disk connectivity, because the unaffected plex continues to serve data while you fix the cause of the failure. Once the plex that had a problem is fixed, you can resynchronize the two plexes and reestablish the mirror relationship. For more information about snapshots, SnapMirror, and SyncMirror, see the *Data Protection Online Backup and Recovery Guide*.

In the following diagram, SyncMirror is enabled and implemented, so plex0 has been copied and automatically named *plex1* by Data ONTAP. Plex0 and plex1 contain copies of one or more file systems. In this diagram, thirty-two disks had been available prior to the SyncMirror relationship being initiated. After initiating SyncMirror, each pool has its own collection of sixteen hot spare disks.



When you create an aggregate, Data ONTAP assigns data disks and parity disks to RAID groups, depending on the options you choose, such as the size of the RAID group (based on the number of disks to be assigned to it) or the level of RAID protection.

Choosing the right size and the protection level for a RAID group depends on the kind of data that you intend to store on the disks in that RAID group. For more information about planning the size of RAID groups, see “Size of RAID groups” on page 24 and Chapter 4, “RAID Protection of Data,” on page 105.

Creating aggregates

About creating aggregates

When a single, unmirrored aggregate is first created, all the disks in the single plex must come from the same disk pool.

How Data ONTAP enforces checksum type rules

As mentioned in Chapter 3, Data ONTAP uses the disk's checksum type for RAID parity checksums. You must be aware of a disk's checksum type because Data ONTAP enforces the following rules when creating aggregates or adding disks to existing aggregates (these rules also apply to creating traditional volumes or adding disks to them):

- An aggregate can have only one checksum type, and it applies to the entire aggregate.
- When you create an aggregate:
  - Data ONTAP determines the checksum type of the aggregate, based on the type of disks available.
  - If enough block checksum disks (BCDs) are available, the aggregate uses BCDs.
  - Otherwise, the aggregate uses zoned checksum disks (ZCDs).
  - To use BCDs when you create a new aggregate, you must have at least the same number of block checksum spare disks available that you specify in the `aggr create` command.
- When you add disks to an existing aggregate:
  - You can add a BCD to either a block checksum aggregate or a zoned checksum aggregate.
  - You cannot add a ZCD to a block checksum aggregate.

If you have a system with both BCDs and ZCDs, Data ONTAP attempts to use the BCDs first. For example, if you issue the command to create an aggregate, Data ONTAP checks to see whether there are enough BCDs available.

- If there *are* enough BCDs, Data ONTAP creates a block checksum aggregate.
- If there *are not* enough BCDs, and there *are no* ZCDs available, the command to create an aggregate fails.
- If there *are not* enough BCDs, and there *are* ZCDs available, Data ONTAP checks to see whether there are enough of them to create the aggregate.

- If there *are not* enough ZCDs, Data ONTAP checks to see whether there are enough mixed disks to create the aggregate.
- If there *are* enough mixed disks, Data ONTAP mixes block and zoned checksum disks to create a zoned checksum aggregate.
- If there *are not* enough mixed disks, the command to create an aggregate fails.

Once an aggregate is created on NetApp filers or NearStore appliances, you cannot change the format of a disk. However, on NetApp gFiler gateways, you can convert a disk from one checksum type to the other with the `disk assign -c block | zoned` command. For more information, see the *gFiler Software, Installation, and Management Guide*.

Creating an aggregate

When you create an aggregate, you must provide the following information:

A name for the aggregate: The names must follow these naming conventions:

- Begin with either a letter or an underscore (_)
- Contain only letters, digits, and underscores
- Contain no more than 255 characters

Disks to include in the aggregate: You specify disks by using the `-d` option and their IDs or by the number of disks of a specified size.

All of the disks in an aggregate must follow these rules:

- Disks must be of the same type (FC-AL, ATA, or SCSI).
- Disks must have the same RPM.

If disks with different speeds are present on a NetApp system (for example, both 10,000 RPM and 15,000 RPM disks), Data ONTAP avoids mixing them within one aggregate. By default, Data ONTAP selects disks

- With the same speed when creating an aggregate in response to the following commands:
  - `aggr create`
  - `vol create`
- That match the speed of existing disks in the aggregate that requires expansion or mirroring in response to the following commands:
  - `aggr add`
  - `aggr mirror`
  - `vol add`
  - `vol mirror`

If you use the `-d` option to specify a list of disks for commands that add disks, the operation will fail if the speeds of the disks differ from each other or differ from the speed of disks already included in the aggregate. The commands for which the `-d` option will fail in this case are `aggr create`, `aggr add`, `aggr mirror`, `vol create`, `vol add`, and `vol mirror`. For example, if you enter `aggr create vol4 -d 9b.25 9b.26 9b.27` and two of the disks are of different speeds, the operation fails.

When using the `aggr create` or `vol create` commands, you can use the `-R` *rpm* option to specify the type of disk to used based on speed. You only need to use this option on appliances that have different disks with different speeds. Typical values for *rpm* are: 5400, 7200, 10000, and 15000. The `-R` option cannot be used with the `-d` option.

If you have any question concerning the speed of a disk that you are planning to specify, NetApp recommends that you first use the `sysconfig -r` command to ascertain the speed of the disks that you want to specify.

Note

It is possible to override the RPM check with option `-f`, but NetApp does not recommend this practice because the resulting aggregate may not meet performance expectations.

Data ONTAP periodically checks if adequate spares are available for the filer. In those checks, only disks with matching or higher speeds are considered as adequate spares. However, if a disk fails and a spare with matching speed is not available, Data ONTAP may use a spare with a different (higher or lower) speed for RAID reconstruction.

Note

If an aggregate happens to include disks with different speeds and adequate spares are present, you can use the `disk replace` command to replace mismatched disks. Data ONTAP will use Rapid RAID Recovery to copy such disks to more appropriate replacements.

Note

If you are setting up aggregates on an FAS270c appliance with two internal filer heads or a system licensed for SnapMover, you might have to assign the disks to one of the filers before creating aggregates on those filers. For more information, see “Software-based disk ownership” on page 71.

For information about creating aggregates, see the na_aggr man page.

To create an aggregate, complete the following steps.

Step	Action
1	View a list of the spare disks on your NetApp appliance. These disks are available for you to assign to the aggregate that you want to create. Enter the following command: `aggr status -s` **Result:** The output of `aggr status -s` lists all the spare disks that you can select for inclusion in the aggregate and their capacities.
2	Enter the following command: **`aggr create` *`aggr_name`* `[-f] [-m] [-n] [-t {raid4 | raid_dp} ] [-r` *`raidsize`*`] [-T disk-type][-R rpm]` *`disk-list`*** *aggr_name* is the name for the new aggregate. `-f` overrides the default behavior that does not permit disks in a plex to span disk pools. This option also allows you to mix disks with different RMP speeds. `-m` specifies the optional creation of a SyncMirror-replicated volume if you want to supplement RAID protection with SyncMirror protection. A SyncMirror license is required for this feature. `-t {raid4 | raid_dp}` specifies the type of RAID protection you want to provide for this aggregate. If no RAID type is specified, the default value (`raid_dp`) is applied. `-r` *raidsize* is the maximum number of disks that you want RAID groups created in this aggregate to consist of. If the last RAID group created contains fewer disks than the value specified, any new disks that are added to this aggregate are added to this RAID group until that RAID group reaches the number of disks specified. When that point is reached, a new RAID group will be created for any additional disks added to the aggregate.

Step	Action
	`-T` *disk-type* specifies one of the following types of disk to be used: ATA, FCAL, and SCSI. This option is only needed when creating aggregates on systems that have mixed disks. Mixing disks of different types in one aggregate is not allowed. You cannot use the `-T` option in combination with the `-d` option. `-R` *rpm* specifies the type of disk to used based on its speed. Use this option only on appliances having different disks with different speeds. Typical values for *rpm* are: 5400, 7200, 10000, and 15000. The `-R` option cannot be used with the `-d` option. *disk-list* is one of the following: ◆ *ndisks*[*@disk-size*] *ndisks* is the number of disks to use. It must be at least 2 (3 if RAID-DP is configured). *disk-size* is the disk size to use, in gigabytes. You must have at least *ndisks* available disks of the size you specify. ◆ `-d` *disk_name1 disk_name2... disk_nameN* *disk_name1*, *disk_name2*, and *disk_nameN* are disk IDs of one or more available disks; use a space to separate multiple disks.
3	Enter the following command to verify that the aggregate exists as you specified: **`aggr status aggr_name -r`** *aggr_name* is the name of the aggregate whose existence you want to verify. **Result:** The system displays the RAID groups and disks of the specified aggregate on your filer.

Aggregate creation example: The following command creates an aggregate called newaggr, with no more than eight disks in a RAID group consisting of the disks with disk IDs 8.1, 8.2, 8.3, and 8.4.

```
aggr create newaggr -r 8 -d 8.1 8.2 8.3 8.4.
```

Changing the state of an aggregate

Aggregate states

An aggregate can be in one of the following three states:

- Offline—Read or write access is not allowed (aggregates cannot be taken offline if they still contain FlexVol volumes).
- Restricted—Some operations, such as parity reconstruction, are allowed, but data access is not allowed (aggregates cannot be made restricted if they still contain FlexVol volumes).
- Online—Read and write access to volumes hosted on this aggregate is allowed. An online aggregate can be further described as follows:
  - `copying`—The aggregate is currently the target aggregate of an active `aggr copy` operation.
  - `degraded`—The aggregate contains at least one degraded RAID group that is not being reconstructed after single disk failure.
  - `double degraded`—The aggregate contains at least one RAID group with double disk failure that is not being reconstructed (this state is possible if RAID-DP protection is enabled for the affected aggregate).
  - `double reconstruction xx% complete`—At least one RAID group in the aggregate is being reconstructed after experiencing double disk failure (this state is possible if RAID-DP protection is enabled for the affected aggregate).
  - `foreign`—Disks that the aggregate contains were moved to the current filer from another filer.
  - `growing`—Disks are in the process of being added to the aggregate.
  - `initializing`—The aggregate is in the process of being initialized.
  - `invalid`—The aggregate does not contain a valid file system.
  - `ironing`—A WAFL consistency check is being performed on the aggregate.
  - `mirrored`—The aggregate is mirrored and all of its RAID groups are functional.
  - `mirror degraded`—The aggregate is a mirrored aggregate and one of its plexes is offline or resynchronizing.
  - `needs check`—WAFL consistency check needs to be performed on the aggregate.
  - `normal`—The aggregate is unmirrored and all of its RAID groups are functional.

- `partial`—At least one disk was found for the aggregate, but two or more disks are missing.
- `reconstruction xx% complete`—At least one RAID group in the aggregate is being reconstructed after experiencing single disk failure.
- `resyncing`—The aggregate contains two plexes; one plex is resynchronizing with the aggregate.
- `verifying`—A mirror verification operation is currently running on the aggregate.
- `wafl inconsistent`—The aggregate has been marked corrupted; contact NetApp Technical Support.

Determining the state of aggregates

To determine what state an aggregate is in, complete the following step.

Step	Action

1	Enter the following command:
	`aggr status`
	This command displays a concise summary of all the aggregates and traditional volumes in the appliance.
	Example: In the following example, the *State* column displays whether the aggregate is online, offline, or restricted. The *Status* column displays the RAID type and, lists any status other than normal (in the case of volA, below, the status is mirrored).
	`> aggr status`
	`Aggr Type  State      Status    Options`
	`vol0 AGGR  online     raid4     root,`
	`volA TRAD  online     raid_dp`
	`mirrored`

When to take an aggregate offline

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

- To perform maintenance on the aggregate
- To destroy an aggregate

Taking an aggregate offline

There are two ways to take an aggregate offline, depending on whether Data ONTAP is running in normal or maintenance mode. In normal mode, you must first offline and destroy all of the FlexVol volumes in the aggregate. In maintenance mode, the FlexVol volumes are preserved.

To take an aggregate offline while Data ONTAP is running in normal mode, complete the following steps.

Step	Action
1	Ensure that all FlexVol volumes in the aggregate have been taken offline and destroyed.
2	Enter the following command: **`aggr offline`** ***`aggr_name`*** *aggr_name* is the name of the aggregate to be taken offline.

To enter into maintenance mode and take an aggregate offline, complete the following steps.

Step	Action
1	Turn on or reboot the system. When prompted to do so, press Ctrl-C to display the boot menu.
2	Enter the choice for booting in maintenance mode.
3	Enter the following command: **`aggr offline`** ***`aggr_name`*** *aggr_name* is the name of the aggregate to be taken offline.
4	Halt the system to exit maintenance mode by entering the following command: **`halt`**
5	Reboot the system. The system will reboot in normal mode.

Restricting an aggregate

You only restrict an aggregate if you want it to be the target of an aggregate copy operation. For information about the aggregate copy operation, see the *Data Protection Online Backup and Recovery Guide*.

To restrict an aggregate, complete the following step.

Step	Action
1	Enter the following command: **`aggr restrict aggr_name`** *aggr_name* is the name of the aggregate to be made restricted.

Bringing an aggregate online

You bring an aggregate online to make it available to the filer after you have taken it offline and are ready to put it back in service.

To bring an aggregate online, complete the following step.

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

Renaming an aggregate

Generally, you might want to rename aggregates to give them descriptive names.

To rename an aggregate, complete the following step.

Step	Action
1	Enter the following command: **`aggr rename aggr_name new_name`** *aggr_name* is the name of the aggregate you want to rename. *new_name* is the new name of the aggregate. **Result:** The aggregate is renamed.

Adding disks to aggregates

Rules for adding disks to an aggregate

You can add disks of various sizes in an aggregate, using the following rules:

- You can add only hot spare disks to an aggregate.
- You must specify the aggregate to which you are adding the disks.
- If you are using mirrored aggregates, the disks must come from the same spare disk pool.
- If the added disk replaces a failed data disk, its capacity is limited to that of the failed disk.
- If the added disk is not replacing a failed data disk and it is not larger than the parity disk, its full capacity (subject to rounding) is available as a data disk.
- If the added disk is larger than an existing parity disk, see "Adding disks larger than the parity disk" on page 168.

If you want to add disks of different speeds, follow the guidelines described in the section about "Disks must have the same RPM." on page 158.

Checksum type rules for creating or expanding aggregates

You must use disks of the appropriate checksum type to create or expand aggregates, as described in the following rules.

- You can add a BCD to a block checksum aggregate or a zoned checksum aggregate.
- You cannot add a ZCD to a block checksum aggregate. For information, see "How Data ONTAP enforces checksum type rules" on page 157.
- To use block checksums when you create a new aggregate, you must have at least the number of block checksum spare disks available that you specified in the `aggr create` command.

The following table shows the types of disks that you can add to an existing aggregate of each type.

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

Hot spare disk planning for aggregates

To fully support an aggregate's RAID disk failure protection, at least one hot spare disk is required for that aggregate. As a result, the filer should contain spare disks of sufficient number and capacity to

- Support the size of the aggregate that you want to create
- Serve as replacement disks should disk failure occur in any aggregate

Note

The size of the spare disks should be equal to or greater than the size of the aggregate disks that the spare disks might replace.

NetApp recommends that you install at least one spare disk matching the size and speed of each aggregate disk.

Adding disks larger than the parity disk

If an added disk is larger than an existing parity disk, the added disk replaces the smaller disk as the parity disk, and the smaller disk becomes a data disk. This enforces a Data ONTAP rule that the parity disk must be at least as large as the largest data disk in a RAID group.

Note

In aggregates configured with RAID-DP, the larger added disk replaces any smaller regular parity disk, but its capacity is reduced, if necessary, to avoid exceeding the capacity of the smaller-sized dParity disk.

Adding disks to an aggregate

To add new disks to an aggregate or a traditional volume, complete the following steps.

Step	Action
1	Verify that hot spare disks are available for you to add by entering the following command: `aggr status -s`

Step	Action
2	Add the disks by entering the following command: **`aggr add aggr_name [-f] [-n] {ndisks[@disk-size]	[-d disk1 [disk2 ...] [disk1 [disk2 ...]] }`** *aggr_name* is the name of the aggregate to which you are adding the disks. `-f` overrides the default behavior that does not permit disks in a plex to span disk pools (only applicable if SyncMirror is licensed). This option also allows you to mix disks with different speeds. `-n` causes the command that Data ONTAP will execute to be displayed without actually executing the command. This is useful for displaying the disks that would be automatically selected prior to executing the command. *ndisks* is the number of disks to use. *disk-size* is the disk size, in gigabytes, to use. You must have at least *ndisks* available disks of the size you specify. `-d` specifies that the *disk-name* will follow. If the aggregate is mirrored, then the `-d` argument must be used twice (if you are specifying *disk-names*). *disk-name* is the disk number of a spare disk; use a space to separate disk numbers. The disk number is under the Device column in the `aggr status -s` display. **Note** If you want to use block checksum disks in a zoned checksum aggregate even though there are still zoned checksum hot spare disks, use the `-d` option to select the disks. **Examples:** The following command adds four 72-GB disks to the thisaggr aggregate: `aggr add thisaggr 4@72` The following command adds the disks 7.17 and 7.26 to the thisaggr aggregate: `aggr add thisaggr -d 7.17 7.26`

Adding disks to a specific RAID group in an aggregate

If an aggregate has more than one RAID group, you can specify the RAID group to which you are adding disks. To add disks to a specific RAID group of an aggregate, complete the following step.

Step	Action
1	Enter the following command: `aggr add` *`aggr_name`* `-g` *`raidgroup ndisks`*`[@`*`disk-size`*`] | -d` *`disk-name...`* *raidgroup* is a RAID group in the aggregate specified by *aggr_name* **Example:** The following command adds two disks to RAID group 0 of the vol0 volume: `aggr add aggr0 -g rg0 2`

The number of disks you can add to a specific RAID group is limited by the `raidsize` setting of the aggregate to which that group belongs. For more information, see Chapter 4, “Changing the size of existing RAID groups,” on page 129

Forcibly adding disks to aggregates

If you try to add disks to an aggregate (or traditional volume) under the following situations, the operation will fail:

- The disks specified in the `aggr add` (or `vol add`) command would cause the disks on a mirrored aggregate to consist of disks from two spare pools.
- The disks specified in the `aggr add` (or `vol add`) command have a different speed in revolutions per minute (RPM) than that of existing disks in the aggregate.

If you add disks to an aggregate (or traditional volume) under the following situation, the operation will prompt you for confirmation, and then succeed or abort, depending on your response.

- The disks specified in the `aggr add` command would add disks to a RAID group other than the last RAID group, thereby making it impossible for the file system to revert to an earlier version than Data ONTAP 6.2.

To force Data ONTAP to add disks in these situations, complete the following step.

Step	Action
1	Enter the following command: `aggr add aggr-name -f [-g raidgroup] -d disk-name...` **Note** You must use the `-g` *raidgroup* option to specify a RAID group other than the last RAID group in the aggregate.

Displaying disk space on an aggregate

To display the disk space available on an aggregate, complete the following step.

Step	Action

1	Enter the following command:
	`df -A aggr_name`
	Example:
	`df -A aggr1`
	`Aggregate kbytes used avail capacity`
	`aggr1 4339168 1777824 2561344 41%`
	`aggr1/.snapshot 1084788 956716 128072 88%`

After adding disks for LUNs, you run reallocation jobs

After you add disks to an aggregate, NetApp recommends that you run a full reallocation job on each FlexVol contained in that aggregate. For information on how to perform this task, see the *Block Access Management Guide for FCP* or the *Block Access Management Guide for iSCSI*.

Destroying aggregates

About destroying aggregates

You cannot destroy an aggregate while it still contains any FlexVol volumes.

There are two reasons to destroy an aggregate:

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

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

Caution

If you destroy an aggregate, all the data in the aggregate is destroyed and no longer accessible.

Note

You can destroy a SnapLock Enterprise aggregate at any time; however, you cannot destroy a SnapLock Compliance aggregate until the retention periods for all data contained in it have expired.

Destroying an aggregate

To destroy an aggregate, complete the following steps.

Step	Action

1	Take all FlexVol volumes offline and destroy them by entering the following commands for each volume:
	`vol offline vol_name`
	`vol destroy vol_name`

2	Take the aggregate offline by entering the following command: `aggr offline aggr_name` *aggr_name* is the name of the aggregate that you intend to destroy and whose disks you are converting to hot spares.

Step	Action
3	Destroy the aggregate by entering the following command: **`aggr destroy aggr_name`** *aggr_name* is the name of the aggregate that you intend to destroy and whose disks you are converting to hot spares.

Physically moving aggregates

About physically moving aggregates

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

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

You can physically move disks, disk shelves, or loops to move an aggregate from one filer to another.

The following terms are used:

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

NetApp recommends that you only move disks from a source filer to a destination filer if the destination filer has higher NVRAM capacity.

Note

The procedure described here does not apply to gFilers. For information about how to physically move aggregates in gFilers, see the *gFiler Software Setup, Installation, and Management Guide*.

Physically moving an aggregate

To physically move an aggregate, complete the following steps.

Step	Action
1	In normal mode, enter the following command at the source filer to locate the disks that contain the aggregate: `aggr status aggr_name -r` **Result:** The locations of the data and parity disks in the aggregate appear under the aggregate name on the same line as the labels Data and Parity.
2	Reboot the source filer into maintenance mode.

Step	Action
3	In maintenance mode, take the aggregate that you want to move offline. **`aggr offline aggr_name`** Then follow the instructions in the disk shelf hardware guide to remove the disks from the source filer.
4	Halt and turn off the destination filer.
5	Install the disks in a disk shelf connected to the destination filer.
6	Reboot the destination filer in maintenance mode. **Result:** When the destination filer boots, it takes the foreign aggregate offline. If the foreign aggregate has the same name as an existing aggregate on the filer, the filer renames it *aggr_name(1)*, where *aggr_name* is the original name of the aggregate. **Caution** If the foreign aggregate is incomplete, repeat Step 5 to add the missing disks. Do not try to add missing disks while the aggregate is online—doing so causes them to become hot spare disks.
7	If the filer renamed the foreign aggregate because of a name conflict, enter the following command to rename the aggregate: **`aggr rename aggr_name new_name`** *aggr_name* is the name of the aggregate you want to rename. *new_name* is the new name of the aggregate. **Example:** The following command renames the users(1) aggregate as newusers: `aggr rename users(1) newusers`

Step	Action
8	Enter the following command to bring the aggregate online in the destination filer: **`aggr online aggr_name`** *aggr_name* is the name of the aggregate. **Result:** The aggregate is online in its new location in the destination filer.
9	Enter the following command to confirm that the added aggregate came online: **`aggr status aggr_name`** *aggr_name* is the name of the aggregate.
10	Power up and reboot the source filer.
11	Reboot the destination filer out of maintenance mode.

Volume Management 6

About this chapter

This chapter describes how to use volumes to contain and manage user data.

Topics in this chapter

This chapter discusses the following topics:

- "Traditional and FlexVol volumes" on page 178
- "Traditional volume operations" on page 181
- "FlexVol operations" on page 190
- "General volume operations" on page 204
- "Managing FlexCache volumes" on page 229
- "Space management for volumes and files" on page 244

Traditional and FlexVol volumes

About traditional and FlexVol volumes

Volumes are file systems that hold user data that is accessible via one or more of the access protocols supported by Data ONTAP, including NFS, CIFS, HTTP, WebDAV, FTP, FCP and iSCSI. You can create one or more snapshots of the data in a volume so that multiple, space-efficient, point-in-time images of the data can be maintained for such purposes as backup and error recovery.

Each volume depends on its containing aggregate for all its physical storage, that is, for all storage in the aggregate's disks and RAID groups. A volume is associated with its containing aggregate in one of the two following ways:

- A *traditional volume* is a volume that is contained by a single, dedicated, aggregate; it is tightly coupled with its containing aggregate. The only way to grow a traditional volume is to add entire disks to its containing aggregate. It is impossible to decrease the size of a traditional volume. The smallest possible traditional volume must occupy all of two disks (for RAID-4) or three disks (for RAID-DP).

  No other volumes can get their storage from this containing aggregate.

  All volumes created with a version of Data ONTAP earlier than 7.0 are traditional volumes. If you upgrade to Data ONTAP 7.0 or later, your volumes and data remain unchanged, and the commands you used to manage your volumes and data are still supported.
- A *FlexVol* (sometimes called a flexible volume) is a volume that is loosely coupled to its containing aggregate. Because the volume is managed separately from the aggregate, you can create small FlexVol volumes (20 MB or larger), and you can increase or decrease the size of FlexVol volumes in increments as small as 4 KB.

  A FlexVol can share its containing aggregate with other FlexVol volumes. Thus, a single aggregate can be the shared source of all the storage used by all the FlexVol volumes contained by that aggregate.

Note

FlexVol volumes have different best practices, optimal configurations, and performance characteristics compared to traditional volumes. Make sure you understand these differences and deploy the configuration that is optimal for your environment.

For information about deploying a storage solution with FlexVol volumes, including migration and performance considerations, see the technical report *Introduction to Data ONTAP Release 7G* (available from the NetApp Library at http://www.netapp.com/tech_library/ftp/3356.pdf).

Limits on how many volumes you can have

You can create up to 200 FlexVol and traditional volumes on a single filer. In addition, the following limits apply.

Traditional volumes: You can have up to 100 traditional volumes and aggregates combined on a single filer.

FlexVol volumes: The only limit imposed on FlexVol volumes is the overall system limit of 200 for all volumes.

For clusters, these limits apply to each node individually, so the overall limits for the pair are doubled.

Types of volume operations

The volume operations described in this chapter fall into three types:

- "Traditional volume operations" on page 181

  These are RAID and disk management operations that pertain only to traditional volumes.

  - "Creating traditional volumes" on page 182
  - "Physically transporting traditional volumes" on page 187
- "FlexVol operations" on page 190

  These are operations that use the advantages of FlexVol volumes, so they pertain only to FlexVol volumes.

  - "Creating FlexVol volumes" on page 191
  - "Resizing FlexVol volumes" on page 195
  - "Cloning FlexVol volumes" on page 197
  - "Displaying a FlexVol's containing aggregate" on page 203

- "General volume operations" on page 204

  These are operations that apply to both FlexVol and traditional volumes.

  - "Migrating between traditional volumes and FlexVol volumes" on page 205
  - "Managing volume languages" on page 213
  - "Determining volume status and state" on page 216
  - "Renaming volumes" on page 223
  - "General volume operations" on page 224
  - "Destroying volumes" on page 224
  - "Increasing the maximum number of files in a volume" on page 226
  - "Reallocating file and volume layout" on page 228

Traditional volume operations

About traditional volume operations

Operations that apply exclusively to traditional volumes generally involve management of the disks assigned to those volumes and the RAID groups to which those disks belong.

Traditional volume operations described in this section include:

- "Creating traditional volumes" on page 182
- "Physically transporting traditional volumes" on page 187

Additional traditional volume operations that are described in other chapters or other guides include:

- Configuring and managing RAID protection of volume data

  See "RAID Protection of Data" on page 105.
- Configuring and managing SyncMirror replication of volume data

  See the *Data Protection Online Backup And Recovery Guide*.
- Increasing the size of a traditional volume

  To increase the size of a traditional volume, you increase the size of its containing aggregate. For more information about increasing the size of an aggregate, see "Adding disks to aggregates" on page 167.
- Configuring and managing SnapLock volumes

  See "About SnapLock" on page 328.

Creating traditional volumes

About creating traditional volumes

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

- A name for the volume

  For more information about volume naming conventions, see "Volume naming conventions" on page 182.
- An optional language for the volume

  The default value is the language of the root volume.

  For more information about choosing a volume language, see "Managing volume languages" on page 213.
- The RAID-related parameters for the aggregate that contains the new volume

  For a complete description of RAID related options for volume creation see "Setting RAID type and group size" on page 119.

Volume naming conventions

You choose the volume names. The names must follow these naming conventions:

- Begin with either a letter or an underscore (_)
- Contain only letters, digits, and underscores
- Contain no more than 255 characters

Creating a traditional volume

To create a traditional volume, complete the following steps.

Step	Action
1	At the system prompt, enter the following command: `aggr status -s` **Result:** The output of `aggr status -s` lists all the hot-swappable spare disks that you can assign to the traditional volume and their capacities. **Note** If you are setting up traditional volumes on an FAS270c system with two internal filer heads, or a system that has SnapMover licensed, you might have to assign the disks before creating volumes on those filers. For more information, see “Software-based disk ownership” on page 71.

Step	Action
2	At the system prompt, enter the following command: **`aggr create vol_name -v [-l language_code] [-f] [-n] [-m] [-t raid-type] [-r raid-size] [-T disk-type] [-R rpm] [-L] disk-list`** *vol_name* is the name for the new volume (without the /vol/ prefix). *language_code* specifies the language for the new volume. The default is the language of the root volume. See "Viewing the language list online" on page 214. The `-L` flag is used only when creating SnapLock volumes. For more information about SnapLock volumes, see "SnapLock Management" on page 327. **Note** For a complete description of the all the options for the aggr command, see "Creating an aggregate" on page 158. For information about RAID related options for `aggr create`, see "Setting RAID type and group size" on page 119 or the na_aggr(1) man page. For backward compatibility, you can also use the `vol create` command to create a traditional volume. However, not all of the RAID related options are available for the `vol` command. For more information, see the na_vol(1) man page. **Result:** The new volume is created and, if NFS is in use, an entry for the new volume is added to the /etc/export file. **Example:** The following command creates a traditional volume called newvol, with no more than eight disks in a RAID group, using the French character set, and consisting of the disks with disk IDs 8.1, 8.2, 8.3, and 8.4. `aggr create newvol -v -r 8 -l fr -d 8.1 8.2 8.3 8.4`

Step	Action
3	Enter the following command to verify that the volume exists as you specified: `aggr status vol_name -r` *vol_name* is the name of the volume whose existence you want to verify. **Result:** The system displays the RAID groups and disks of the specified volume on your filer.
4	If you access the filer using CIFS, update your CIFS shares as necessary.
5	If you access the filer using NFS, complete the following steps: 1. Verify that the line added to the /etc/exports file for the new volume is correct for your security model. 2. Add the appropriate mount point information to the /etc/fstab or /etc/vfstab file on clients that mount volumes from the filer.

Parameters to accept or change after volume creation

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

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

For more information about CIFS oplocks, see "Changing the CIFS oplocks setting" on page 265.

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

For more information about file security styles, see "Understanding security styles" on page 261.

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

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

When you change the configuration of a filer from one protocol to another (by licensing or unlicensing protocols), the default security style for new volumes changes as shown in the following table.

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

Checksum type usage

A checksum type applies to an entire traditional volume. A traditional volume can have only one checksum type. For more information about checksum types, see “How Data ONTAP enforces checksum type rules” on page 157.

Physically transporting traditional volumes

About physically moving traditional volumes

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

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

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

The following terms are used:

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

Note

If MultiStore® and SnapMover licenses are installed, you might be able to move traditional volumes without moving the drives on which they are located. For more information, see the *MultiStore Management Guide*.

Moving a traditional volume

To physically move a traditional volume, perform the following steps.

Step	Action
1	Enter the following command at the source filer to locate the disks that contain the volume *vol_name*: `aggr status vol_name -r` **Result:** The locations of the data and parity disks in the volume are displayed.

Step	Action
2	Enter the following command on the source filer to take the volume and its containing aggregate offline: **`aggr offline vol_name`**
3	Follow the instructions in the disk shelf hardware guide to remove the data and parity disks for the specified volume from the source filer.
4	Follow the instructions in the disk shelf hardware guide to install the disks in a disk shelf connected to the destination filer. **Result:** When the destination filer sees the disks, it places the foreign volume offline. If the foreign volume has the same name as an existing volume on the filer, the filer renames it *vol_name(d)*, where *vol_name* is the original name of the volume and *d* is a digit that makes the name unique.
5	Enter the following command to make sure that the newly moved volume is complete: **`aggr status new_vol_name`** *new_vol_name* is the (possibly new) name of the volume you just moved. **Caution** If the foreign volume is incomplete (it has a status of `partial`), add all missing disks before proceeding. Do not try to add missing disks after the volume comes online—doing so causes them to become hot spare disks. You can identify the disks currently used by the volume using the `aggr status -r` command.
6	If the filer renamed the foreign volume because of a name conflict, enter the following command on the target filer to rename the volume: **`aggr rename new_vol_name vol_name`** *new_vol_name* is the name of the volume you want to rename. *vol_name* is the new name of the volume.

Step	Action
7	Enter the following command on the target filer to bring the volume and its containing aggregate online: **`aggr online vol_name`** *vol_name* is the name of the newly moved volume. **Result:** The volume is brought online on the target filer.
8	Enter the following command to confirm that the added volume came online: **`aggr status vol_name`** *vol_name* is the name of the newly moved volume.
9	If you access the filers using CIFS, update your CIFS shares as necessary.
10	If you access the filers using NFS, complete the following steps for both the source and the destination filer: **1.** Update the system /etc/exports file. **2.** Run `exportfs -a`. **3.** Update the appropriate mount point information to the /etc/fstab or /etc/vfstab file on clients that mount volumes from the filer.

FlexVol operations

About FlexVol operations

These operations apply exclusively to FlexVol volumes because they take advantage of the virtual nature of FlexVol volumes.

FlexVol operations described in this section include:

- "Creating FlexVol volumes" on page 191
- "Resizing FlexVol volumes" on page 195
- "Cloning FlexVol volumes" on page 197
- "Displaying a FlexVol's containing aggregate" on page 203

Creating FlexVol volumes

About creating FlexVol volumes

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

- A name for the volume
- The name of the containing aggregate
- The size of the volume

  The size of a FlexVol must be at least 20 MB. The maximum size is 16 TB, or what your filer configuration can support.

In addition, you can provide the following optional values:

- The language used for file names

  The default language is the language of the root volume.
- The space guarantee setting for the new volume

  For more information, see "Space guarantees" on page 247.

Volume naming conventions

You choose the volume names. The names must follow these naming conventions:

- Begin with either a letter or an underscore (_)
- Contain only letters, digits, and underscores
- Contain no more than 255 characters

Creating a FlexVol

To create a FlexVol, complete the following steps.

Step	Action
1	If you have not already done so, create one or more aggregates to contain the FlexVol volumes that you want to create. To view a list of the aggregates that you have already created, and the volumes that they contain, enter the following command: `aggr status -v`

Step	Action
2	At the system prompt, enter the following command: **vol create *f_vol_name* [-l *language_code*] [-s {volume|file|none}] *aggr_name* *size*{k|m|g|t}** *f_vol_name* is the name for the new FlexVol (without the /vol/ prefix). This name must be different from all other volume names on the filer. *language_code* specifies a language other than that of the root volume. See "Viewing the language list online" on page 214. `-s {volume|file|none}` specifies the space guarantee setting that is enabled for the specified FlexVol. If no value is specified, the default value is `volume`. For more information, see "Space guarantees" on page 247. *aggr_name* is the name of the containing aggregate for this FlexVol. *size* { `k|m|g|t` } specifies the volume size in kilobytes, megabytes, gigabytes, or terabytes. For example, you would enter `20m` to indicate twenty megabytes. If you do not specify a unit, *size* is taken as bytes and rounded up to the nearest multiple of 4 KB. **Result:** The new volume is created and, if NFS is in use, an entry is added to the /etc/export file for the new volume. **Example:** The following command creates a 200-MB volume called newvol, in the aggregate called aggr1, using the French character set. `vol create newvol -l fr aggr1 200M`
3	Enter the following command to verify that the volume exists as you specified: **vol status *f_vol_name*** *f_vol_name* is the name of the FlexVol whose existence you want to verify.
4	If you access the filer using CIFS, update the share information for the new volume.

Step	Action
5	If you access the filer using NFS, complete the following steps: 1. Verify that the line added to the /etc/exports file for the new volume is correct for your security model. 2. Add the appropriate mount point information to the /etc/fstab or /etc/vfstab file on clients that mount volumes from the filer.

Parameters to accept or change after volume creation

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

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

For more information about CIFS oplocks, see "Changing the CIFS oplocks setting" on page 265.

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

For more information about file security styles, see "Understanding security styles" on page 261.

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

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

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

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

Resizing FlexVol volumes

About resizing FlexVol volumes

You can increase or decrease the amount of space that an existing FlexVol can occupy on its containing aggregate. A FlexVol can grow to the size you specify as long as the containing aggregate has enough free space to accommodate that growth.

Resizing a FlexVol

To resize a FlexVol, complete the following steps.

Step	Action
1	Check the available space of the containing aggregate by entering the following command: **`df -A aggr_name`** *aggr_name* is the name of the containing aggregate for the FlexVol whose size you want to change.
2	If you want to determine the current size of the volume, enter one of the following commands: **`vol size f_vol_name`** **`df f_vol_name`** *f_vol_name* is the name of the FlexVol that you intend to resize.

Step	Action
3	Enter the following command to resize the volume: **`vol size f_vol_name [+|-] n{k|m|g|t}`** *f_vol_name* is the name of the FlexVol that you intend to resize. If you include the `+` or `-`, *n*{`k`|`m`|`g`|`t`} specifies how many kilobytes, megabytes, gigabytes or terabytes to increase or decrease the volume size. If you do not specify a unit, size is taken as bytes and rounded up to the nearest multiple of 4 KB. If you omit the + or -, the size of the volume is set to the size you specify, in kilobytes, megabytes, gigabytes, or terabytes. If you do not specify a unit, size is taken as bytes and rounded up to the nearest multiple of 4 KB. **Note** If you attempt to decrease the size of a FlexVol to less than the amount of space that it is currently using, the command fails.
4	Verify the success of the resize operation by entering the following command: **`vol size f_vol_name`**

Cloning FlexVol volumes

About cloning FlexVol volumes

Data ONTAP provides the ability to clone FlexVol volumes, creating FlexClone volumes. The following list outlines some key facts about FlexClone volumes that you should know:

- FlexClone volumes are a point-in-time, writable copy of the parent volume. Changes made to the parent volume after the FlexClone is created are not reflected in the FlexClone.
- FlexClone volumes are fully functional volumes; you manage them using the `vol` command, just as you do the parent volume.
- FlexClone volumes always exist in the same aggregate as their parent volumes.
- FlexClone volumes can themselves be cloned.
- FlexClone volumes and their parent volumes share the same disk space for any data common to the clone and parent. This means that creating a FlexClone is instantaneous and requires no additional disk space (until changes are made to the clone or parent).
- Because creating a FlexClone does not involve copying data, FlexClone creation is very fast.
- FlexClone volumes are created with the same space guarantees as their parent.

  Note

  In Data ONTAP 7.0 and later versions, space guarantees are disabled for FlexClone volumes.

  For more information, see "Space guarantees" on page 247.
- While a FlexClone exists, some operations on its parent are not allowed.

  For more information about these restrictions, see "Limitations of volume cloning" on page 199.
- If, at a later time, you decide you want to sever the connection between the parent and the clone, you can split the FlexClone. This removes all restrictions on the parent volume and enables the space guarantee on the FlexClone.

Note

Splitting a FlexClone from its parent volume deletes all existing snapshots of the FlexClone.

For more information, see "About splitting a FlexClone from its parent volume" on page 201.

- When a FlexClone is created, quotas are reset on the FlexClone, and any LUNs present in the parent volume are present in the FlexClone but are unmapped.

  For more information about using volume cloning with LUNs, see the *Block Access Management Guide for iSCSI* or the *Block Access Management Guide for FCP.*
- Only FlexVol volumes can be cloned. To create a copy of a traditional volume, you must use the `vol copy` command, which creates a distinct copy with its own storage.
- You must install the license for the FlexClone feature before you can create FlexClone volumes.

Uses of volume cloning

You can use volume cloning whenever you need a writable, point-in-time copy of an existing FlexVol, including the following scenarios:

- You need to create a temporary copy of a volume for testing purposes.
- You need to make a copy of your data available to additional users without giving them access to the production data.
- You want to create a clone of a database for manipulation and projection operations, while preserving the original data in unaltered form.

Benefits of volume cloning versus volume copying

Volume cloning provides similar results to volume copying, but cloning offers some important advantages over volume copying:

- Volume cloning is instantaneous, whereas volume copying can be time consuming.
- If the original and cloned volumes share a large amount of identical data, considerable space is saved because the shared data is not duplicated between the volume and the clone.

Limitations of volume cloning

The following operations are not allowed on parent volumes or their clones.

- You cannot delete the base snapshot of a parent volume while a cloned volume exists.

  The base snapshot is the snapshot that was used to create the FlexClone. It is named clone_*clonename*.*x*, where *clonename* is the name of the FlexClone and *x* is a numeral that distinguishes it from other snapshots.
- You cannot perform a volume SnapRestore® operation on the parent volume using a snapshot that was taken before the base snapshot was taken.
- You cannot destroy a parent volume if any clone of that volume exists.
- You cannot clone a volume that has bee taken offline, although you can take the parent volume offline after it has been cloned.
- You cannot use either the parent volume or the clone as a destination for a `vol copy` command or volume SnapMirror.
- In Data ONTAP 7.0 and later versions, space guarantees are disabled for FlexClone volumes. This means that writes to a FlexClone can fail if its containing aggregate does not have enough available space, even for LUNs or files with space reservations enabled.

Cloning a FlexVol

To create a FlexClone by cloning a FlexVol, complete the following steps.

Step	Action
1	Ensure that you have the flex_clone license installed.

Step	Action
2	Enter the following command to clone the volume: **`vol clone create cl_vol_name [-s {volume|file|none}] -b f_p_vol_name [parent_snap]`** *cl_vol_name* is the name of the FlexClone that you want to create. `-s {volume | file | none}` specifies the space guarantee setting for the new FlexClone. If no value is specified, the FlexClone is given the same space guarantee setting as its parent. For more information, see "Space guarantees" on page 247. **Note** For Data ONTAP 7.0 and later versions, space guarantees are disabled for FlexClone volumes. *f_p_vol_name* is the name of the FlexVol that you intend to clone. *parent_snap* is the name of the base snapshot of the parent FlexVol. If no name is specified, Data ONTAP creates a base snapshot with the name clone_*cl_name_prefix.id*, where *cl_name_prefix* is the name of the new FlexClone (up to 16 characters) and *id* is a unique digit identifier (for example 1,2, etc.). The base snapshot cannot be deleted as long as the parent volume or any of its clones exists. **Result:** The FlexClone is created and, if NFS is in use, an entry is added to the /etc/export file for every entry found for the parent volume. **Example snapshot name:** To create a FlexClone "newclone" from the parent "flexvol1", the following command is entered: `vol clone create newclone -b flexvol1` The snapshot created by Data ONTAP is named "clone_newclone.1".
3	Verify the success of the FlexClone creation by entering the following command: **`vol status -v cl_vol_name`**

About splitting a FlexClone from its parent volume

You might want to split your FlexClone and its parent into two independent volumes that occupy their own disk space.

Caution

When you split a FlexClone from its parent, all existing snapshots of the FlexClone are deleted.

Because the clone-splitting operation is a copy operation that might take considerable time to carry out, Data ONTAP also provides commands to stop or check the status of a clone-splitting operation.

The clone-splitting operation proceeds in the background and does not interfere with data access to either the parent or the clone volume.

If you take the FlexClone offline while the splitting operation is in progress, the operation is suspended; when you bring the FlexClone back online, the splitting operation resumes.

Once a FlexClone and its parent volume have been split, they cannot be rejoined.

Splitting a FlexClone

To split a FlexClone from its parent volume, complete the following steps.

Step	Action
1	Verify that enough additional disk space exists in the containing aggregate to support storing the data of both the FlexClone and its parent volume, once they are no longer sharing their shared disk space, by entering the following command: **`df -A aggr_name`** *aggr_name* is the name of the containing aggregate of the FlexClone that you want to split. The `avail` column tells you how much available space you have in your aggregate. **Note** When a FlexClone is split from its parent, the resulting two FlexVol volumes occupy completely different blocks within the same aggregate.

Step	Action
2	Enter the following command to split the volume: **`vol clone split start cl_vol_name`** *cl_vol_name* is the name of the FlexClone that you want to split from its parent. **Result:** The original volume and its clone begin to split apart, no longer sharing the blocks that they formerly shared. All existing snapshots of the FlexClone are deleted.
3	If you want to check the status of a clone-splitting operation, enter the following command: **`vol clone status cl_vol_name`**
4	If you want to stop the progress of an ongoing clone-splitting operation, enter the following command: **`vol clone stop cl_vol_name`** **Result:** The clone-splitting operation halts; the original and FlexClone volumes remain clone partners, but the disk space that was duplicated up to that point remains duplicated. All existing snapshots of the FlexClone are deleted.
5	To display status for the newly split FlexVol and verify the success of the clone-splitting operation, enter the following command: **`vol status -v cl_vol_name`**

Displaying a FlexVol's containing aggregate

Showing a FlexVol's containing aggregate

To display the name of a FlexVol's containing aggregate, complete the following step.

Step	Action
1	Enter the following command: **`vol container vol_name`** *vol_name* is the name of the volume whose containing aggregate you want to display.

General volume operations

About general volume operations

General volume operations apply to both traditional volumes and FlexVol volumes.

General volume operations described in this section include:

- "Migrating between traditional volumes and FlexVol volumes" on page 205
- "Managing duplicate volume names" on page 212
- "Managing volume languages" on page 213
- "Determining volume status and state" on page 216
- "Renaming volumes" on page 223
- "Destroying volumes" on page 224
- "Increasing the maximum number of files in a volume" on page 226
- "Reallocating file and volume layout" on page 228

Additional general volume operations that are described in other chapters or other guides include:

- Making a volume available

  For more information on making volumes available to users who are attempting access through NFS, CIFS, FTP, WebDAV, or HTTP protocols, see the *File Access and Protocols Management Guide*.

- Copying volumes

  For more information about copying volumes see the *Data Protection Online Backup and Recovery Guide*.

- Changing the root volume

  For more information about changing the root volume from one volume to another, see the section on the root volume in the *System Administration Guide*.

Migrating between traditional volumes and FlexVol volumes

About migrating between traditional and FlexVol volumes

FlexVol volumes have different best practices, optimal configurations, and performance characteristics compared to traditional volumes. Make sure you understand these differences by referring to the available documentation on FlexVol volumes and deploy the configuration that is optimal for your environment.

For information about deploying a storage solution with FlexVol volumes, including migration and performance considerations, see the technical report *Introduction to Data ONTAP Release 7G* (available from the NetApp Library at http://www.netapp.com/tech_library/ftp/3356.pdf). For information about configuring FlexVol volumes, see "FlexVol operations" on page 190. For information about configuring aggregates, see "Aggregate Management" on page 153.

The following list outlines some facts about migrating between traditional and FlexVol volumes that you should know:

- You cannot convert directly from a traditional volume to a FlexVol, or from a FlexVol to a traditional volume. You must create a new volume of the desired type and then move the data to the new volume using `ndmpcopy`.
- If you move the data to another volume on the same filer, remember that this requires the system to have enough storage to contain both copies of the volume.
- Snapshots on the original volume are unaffected by the migration, but they are not valid for the new volume.

NetApp offers assistance

NetApp Professional Services staff, including Professional Services Engineers (PSEs) and Professional Services Consultants (PSCs) are trained to assist customers with converting volume types and migrating data, among other services. For more information, contact your local NetApp Sales representative, PSE, or PSC.

Migrating a traditional volume to a FlexVol

The following procedure describes how to migrate from a traditional volume to a FlexVol. If you are migrating your root volume, you can use the same procedure, including the steps that are specific to migrating a root volume.

To migrate a traditional volume to a FlexVol, complete the following steps.

Step	Action

1	Determine the size requirements for the new FlexVol. Enter the following command to determine the amount of space your current volume uses:
	df -Ah [*vol_name*]
	Example: `df -Ah vol0`
	Result: The following output is displayed.
	`Aggregate total used avail capacity`
	`vol0 24GB 1434MB 22GB 7%`
	`vol0/.snapshot 6220MB 4864MB 6215MB 0%`
	Root volumes: While it is possible to have small FlexVol volumes, NetApp recommends that root volumes be at least 90 GB so there is enough space for snapshots, log files and coredumps.

2	You can use an existing aggregate or you can create a new one to contain the new FlexVol. To determine if an existing aggregate is large enough to contain the new FlexVol, enter the following command: **df -Ah** **Result:** All of the existing aggregates are displayed.
3	If needed, create a new aggregate by entering the following command: **aggr create *aggr_name disk-list*** **Example:** `aggr create aggrA 8@144` **Result:** An aggregate called *aggrA* is created with eight 144-GB disks. The default RAID type is RAID-DP, so two disks will be used for parity (one parity disk and one dParity disk). The aggregate size will be 1,128 GB. If you want to use RAID-4, and use one less parity disk, enter the following command: **aggr create aggrA -t raid4 8@144**

Step	Action
4	If you want to use the new FlexVol to have the same name as the old traditional volume, you must rename the existing traditional root volume before creating the new FlexVol. Do this by entering the following command: **aggr rename *vol_name new_vol_name*** **Example:** `aggr rename vol0 vol0trad`
5	Create the new FlexVol in the containing aggregate. For more information about creating FlexVol volumes, see "Creating FlexVol volumes" on page 191. **vol create *vol_name aggr_name*** **[-s { volume | file | none }] *size*** **Example:** `vol create vol0 aggrA 90g` **Root volume:** NetApp recommends that you use the (default) `volume` space guarantee for root volumes, because it ensures that writes to the volume do not fail due to a lack of available space in the containing aggregate.
6	Confirm the size of the new FlexVol by entering the following command: **df -h *vol_name***
7	Shut down any applications that use the data to be migrated. Make sure that all data is unavailable to clients and that all files to be migrated are closed.
8	Enable the `ndmpd.enable` option by entering the following command: **options ndmpd.enable on**
9	Migrate the data by entering the following command: **ndmpcopy *old_vol_name new_vol_name*** **Example:** `ndmpcopy /vol/vol0trad /vol/vol0` For more information about using `ndmpcopy`, see the *Data Protection Tape Backup and Recovery Guide*.

Step	Action
10	Verify that the `ndmpcopy` operation completed successfully by verifying that the data was replicated correctly.
11	If you are migrating your root volume, make the new FlexVol the root volume by entering the following command: **`vol options vol_name root`** **Example:** `vol options vol0 root`
12	Reboot the NetApp system.
13	Update the clients to point to the new FlexVol. In a CIFS environment, follow these steps: **1.** Point CIFS shares to the new FlexVol. **2.** Update the CIFS maps on the client machines so that they point to the new FlexVol. In an NFS environment, follow these steps: **1.** Point NFS exports to the new FlexVol. **2.** Update the NFS mounts on the client machines so that they point to the new FlexVol.
14	Make sure all clients can see the new FlexVol and read and write data. To test whether data can be written, complete the following steps: **1.** Create a new folder. **2.** Verify that the new folder exists. **3.** Delete the new folder.
15	If quotas were used with the traditional volume, configure the quotas on the new FlexVol.
16	Take a snapshot of the target volume and create a new snapshot schedule as needed. For more information about taking snapshots, see the *Data Protection Online Backup and Recovery Guide*.

Step	Action
17	When you are confident the volume migration was successful, you can take the original volume offline or destroy it. **Caution** NetApp recommends that you preserve the original volume and its snapshots until the new FlexVol has been stable for some time.

Converting a FlexVol to a traditional volume

To convert a FlexVol to a traditional volume, complete the following steps

Step	Action

1	Determine the size requirements for the new traditional volume. Enter the following command to determine the amount of space your current volume uses:
	df -Ah [*vol_name*]
	Example: df -Ah vol_users
	Result: The following output is displayed.
	`Aggregate total used avail capacity`
	`users 94GB 1434GB 22GB 6%`
	`users/.snapshot 76220MB 74864MB 6215MB 0%`

2	Create the traditional volume that will replace the FlexVol by entering the following command: **aggr create *vol_name disk-list*** **Example: aggr create users 3@144**
3	Confirm the size of the new traditional volume by entering the following command: **df -h *vol_name***
4	Shut down the applications that use the data to be migrated. Make sure that all data is unavailable to clients and that all files to be migrated are closed.

Step	Action
5	Enable the `ndmpd.enable` option by entering the following command: **`options ndmpd.enable on`**
6	Migrate the data using the `ndmpcopy` command. For more information about using `ndmpcopy`, see the *Data Protection Tape Backup and Recovery Guide*.
7	Verify that the `ndmpcopy` operation completed successfully by checking that the data has been replicated correctly.
8	Update the clients to point to the new volume. In a CIFS environment, follow these steps: **1.** Point CIFS shares to the new volume. **2.** Update the CIFS maps on the client machines so that they point to the new volume. **3.** Repeat steps 1 and 2 for each new volume. In an NFS environment, follow these steps: **1.** Point NFS exports to the new volume. **2.** Update the NFS mounts on the client machines so that they point to the new volume. **3.** Repeat steps 1 and 2 for each new volume.
9	Make sure all clients can see the new traditional volume and read and write data. To test whether data can be written, complete the following steps: **1.** Create a new folder. **2.** Verify that the new folder exists. **3.** Delete the new folder. **4.** Repeat steps 1 through 3 for each new volume.
10	If quotas were used with the FlexVol, configure the quotas on the new volume.

Step	Action
11	Take a snapshot of the target volume and create a new snapshot schedule as needed. For more information about taking snapshots, see the *Data Protection Online Backup and Recovery Guide*.
12	When you are confident the volume migration was successful, you can take the source volume offline or destroy it. **Caution** NetApp recommends that you preserve the original volume and its snapshots until the new volume has been stable for some time.

Managing duplicate volume names

How duplicate volume names can occur

Data ONTAP does not support having two volumes with the same name on the same filer. However, certain events can cause this to happen, as outlined in the following list:

- You copy an aggregate using the `aggr copy` command, and when you bring the target aggregate online, there are one or more volumes on the destination filer with the duplicated names.
- You move an aggregate from one filer to another by moving its associated disks, and there is another volume on the destination filer with the same name as a volume contained by the aggregate you moved.
- You move a traditional volume from one filer to another by moving its associated disks, and there is another volume on the destination filer with the same name.
- Using SnapMover, you migrate a vFiler that contains a volume with the same name as a volume on the destination filer.

How Data ONTAP handles duplicate volume names

When Data ONTAP senses a potential duplicate volume name, it appends the string "(*d*)" to the end of the name of the new volume, where *d* is a digit that makes the name unique.

For example, if you have a volume named vol1, and you copy a volume named vol1 from another filer, the newly copied volume might be named vol1(1).

Duplicate volumes should be renamed as soon as possible

You might consider a volume name such as vol1(1) to be acceptable. However, it is important that you rename any volume with an appended digit as soon as possible, for the following reasons:

- The name containing the appended digit is not guaranteed to persist across reboots. Renaming the volume will prevent the name of the volume from changing unexpectedly later on.
- The parentheses characters, "(" and ")", are not legal characters for NFS. Any volume whose name contains those characters cannot be exported to NFS clients.
- The parentheses characters could cause problems for client scripts.

Managing volume languages

About volumes and languages

Every volume has a language. The filer uses a character set appropriate to the language for the following items on that volume:

- File names
- File access

The language of the root volume is used for the following items:

- System name
- CIFS share names
- NFS user and group names
- CIFS user account names
- Domain name
- Console commands and command output
- Access from CIFS clients that don't support Unicode
- Reading the following files:
  - /etc/quotas
  - /etc/usermap.cfg
  - the home directory definition file

Caution

NetApp strongly recommends that all volumes have the same language as the root volume, and that you set the volume language at volume creation time. Changing the language of an existing volume can cause some files to become inaccessible.

Note

Names of the following objects must be in ASCII characters:

- Qtrees
- Snapshots
- Volumes

Viewing the language list online

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

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

Choosing a language for a volume

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

Step	Action	
1	**If the volume is accessed using...**	**Then...**
	NFS Classic (v2 or v3) only	Do nothing; the language does not matter.
	NFS Classic (v2 or v3) and CIFS	Set the language of the volume to the language of the clients.
	NFS v4, with or without CIFS	Set the language of the volume to *cl_lang*.UTF-8, where *cl_lang* is the language of the clients. **Note** If you use NFS v4, all NFS Classic clients must be configured to present file names using UTF-8.

Displaying volume language use

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

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

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

Step	Action

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

Changing the language for a volume

Before changing the language that a volume uses, be sure you read and understand the section titled "About volumes and languages" on page 213.

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

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

Determining volume status and state

Volume states

A volume can be in one of the following three states, sometimes called mount states:

- `online`—Read and write access is allowed.
- `offline`—Read or write access is not allowed.
- `restricted`—Some operations, such as copying volumes and parity reconstruction, are allowed, but data access is not allowed.

Volume status

A volume can have one or more of the following statuses:

Note

Although FlexVol volumes do not directly involve RAID, the state of a FlexVol includes the state of its containing aggregate. Thus, the states pertaining to RAID apply to FlexVol volumes as well as traditional volumes.

- `copying`

  The volume is currently the target volume of active `vol copy` or `snapmirror` operations.

- `degraded`

  The volume's containing aggregate has at least one degraded RAID group that is not being reconstructed.

- `flex`

  The volume is a FlexVol.

- `flexcache`

  The volume is a FlexCache volume. For more information about FlexCache volumes, see "Managing FlexCache volumes" on page 229.

- `foreign`

  Disks used by the volume's containing aggregate were moved to the current filer from another filer.

- `growing`

  Disks are in the process of being added to the volume's containing aggregate.

- `initializing`

  The volume or its containing aggregate are in the process of being initialized.
- `invalid`

  The volume does not contain a valid file system. This typically happens only after an aborted `vol copy` operation.
- `ironing`

  A WAFL consistency check is being performed on the volume's containing aggregate.
- `mirror degraded`

  The volume's containing aggregate is a mirrored aggregate, and one of its plexes is offline or resyncing.
- `mirrored`

  The volume's containing aggregate is mirrored and all of its RAID groups are functional.
- `needs check`

  A WAFL consistency check needs to be performed on the volume's containing aggregate.
- `out-of-date`

  The volume's containing aggregate is mirrored and needs to be resynchronized.
- `partial`

  At least one disk was found for the volume's containing aggregate, but two or more disks are missing.
- `raid0`

  The volume's containing aggregate consists of RAID-0 (no parity) RAID groups (gFiler and NetCache® systems only).
- `raid4`

  The volume's containing aggregate consists of RAID-4 RAID groups.
- `raid_dp`

  The volume's containing aggregate consists of RAID-DP (Double Parity) RAID groups.
- `reconstruct`

  At least one RAID group in the volume's containing aggregate is being reconstructed.

- `resyncing`

  One of the plexes of the volume's containing mirrored aggregate is being resynchronized.
- `snapmirrored`

  The volume is in a SnapMirror relationship with another volume.
- `trad`

  The volume is a traditional volume.
- `unrecoverable`

  The volume is a FlexVol that has been marked unrecoverable. If a volume appears in this state, contact NetApp technical support.
- `verifying`

  A RAID mirror verification operation is currently being run on the volume's containing aggregate.
- `wafl inconsistent`

  The volume or its containing aggregate has been marked corrupted. If a volume appears in this state, contact NetApp technical support.

Determining the state and status of volumes

To determine what state a volume is in, and what status currently applies to it, complete the following step.

Step	Action

1	Enter the following command:
	`vol status`
	This command displays a concise summary of all the volumes in the storage appliance.
	Result: The *State* column displays whether the volume is online, offline, or restricted. The *Status* column displays the volume's RAID type, whether the volume is a FlexVol or traditional volume, and any status other than normal (such as partial or degraded).
	Example:
	`> vol status`
	`Volume State Status Options`
	`vol0 online raid4, flex root,guarantee=volume`
	`volA online raid_dp, trad`
	`mirrored`
	Note
	To see a complete list of all options, including any that are off or not set for this volume, use the `-v` flag with the `vol status` command.

When to take a volume offline

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

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

Note

You cannot take the root volume offline.

Taking a volume offline

To take a volume offline, complete the following step.

Step	Action
1	Enter the following command: `vol offline vol_name` *vol_name* is the name of the volume to be taken offline. **Note** When you take a FlexVol offline, it relinquishes any unused space that has been allocated for it in its containing aggregate. If this space is allocated for another volume and then you bring the volume back online, this can result in an overcommitted aggregate. For more information, see "Bringing a volume online in an overcommitted aggregate" on page 251.

When to make a volume restricted

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

- To copy a volume to another volume

  For more information about volume copy, see the *Data Protection Online Backup and Recovery Guide*.
- To perform a level-0 SnapMirror operation

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

Restricting a volume

To restrict a volume, complete the following step.

Step	Action
1	Enter the following command: `vol restrict vol_name` *vol_name* is the name of the volume to restrict. **Note** When you restrict a FlexVol, it relinquishes any unused space that has been allocated for it in its containing aggregate. If this space is allocated for another volume and then you bring the volume back online, this can result in an overcommitted aggregate. For more information, see "Bringing a volume online in an overcommitted aggregate" on page 251.

Bringing a volume online

You bring a volume back online to make it available to the filer after you deactivated that volume.

Note

If you bring a FlexVol online into an aggregate that does not have sufficient free space in the aggregate to fulfill the space guarantee for that volume, this command fails.

For more information, see "Bringing a volume online in an overcommitted aggregate" on page 251.

To bring a volume back online, complete the following step.

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

Renaming volumes

Renaming a volume

To rename a volume, complete the following steps.

Step	Action
1	Enter the following command: **`vol rename vol_name new-name`** *vol_name* is the name of the volume you want to rename. *new-name* is the new name of the volume. **Result:** The following events occur: ◆ The volume is renamed. ◆ If NFS is in use and the `nfs.exports.auto-update` option is On, the /etc/exports file is updated to reflect the new volume name. ◆ If CIFS is running, shares that refer to the volume are updated to reflect the new volume name. ◆ The in-memory information about active exports gets updated automatically, and clients continue to access the exports without problems.
2	If you access the filer using NFS, add the appropriate mount point information to the /etc/fstab or /etc/vfstab file on clients that mount volumes from the filer.

Destroying volumes

About destroying volumes

There are two reasons to destroy a volume:

- You no longer need the data it contains.
- You copied the data it contains elsewhere.

When you destroy a traditional volume: You also destroy the traditional volume's dedicated containing aggregate. This converts its parity disk and all its data disks back into hot spares. You can then use them in other aggregates, traditional volumes, or filers.

When you destroy a FlexVol: All the disks included in its containing aggregate remain assigned to that containing aggregate.

Caution

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

Destroying a volume

To destroy a volume, complete the following steps.

Step	Action
1	Take the volume offline by entering the following command: `vol offline vol_name` *vol_name* is the name of the volume that you intend to destroy.

Step	Action
2	Enter the following command to destroy the volume: **`vol destroy vol_name`** *vol_name* is the name of the volume that you intend to destroy. **Result:** The following events occur: ◆ The volume is destroyed. ◆ If NFS is in use and the `nfs.exports.auto-update` option is On, entries in the /etc/exports file that refer to the destroyed volume are removed. ◆ If CIFS is running, any shares that refer to the destroyed volume are deleted. ◆ If the destroyed volume was a FlexVol, its allocated space is freed, becoming available for allocation to other FlexVol volumes contained by the same aggregate. ◆ If the destroyed volume was a traditional volume, the disks it used become hot-swappable spare disks.
3	If you access your filer using NFS, update the appropriate mount point information in the /etc/fstab or /etc/vfstab file on clients that mount volumes from the filer.

Increasing the maximum number of files in a volume

About increasing the maximum number of files

The filer automatically sets the maximum number of files for a newly created volume based on the amount of disk space in the volume. The filer increases the maximum number of files when you add a disk to a volume. The number set by the filer never exceeds 33,554,432 unless you set a higher number with the `maxfiles` command. This prevents a filer with terabytes of storage from creating a larger than necessary inode file.

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

Caution

Use caution when increasing the maximum number of files, because after you increase this number, you can never decrease it. As new files are created, the file system consumes the additional disk space required to hold the inodes for the additional files; there is no way for the filer to release that disk space. An inode is a data structure containing information about files.

Increasing the maximum number of files allowed on a volume

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

Step	Action
1	Enter the following command: **`maxfiles vol_name max`** *vol_name* is the volume whose maximum number of files you are increasing. *max* is the maximum number of files. **Note** Inodes are added in blocks, and 5 percent of the total number of inodes is reserved for internal use. If the requested increase in the number of files is too small to require a full inode block to be added, the `maxfiles` value is not increased. If this happens, repeat the command with a larger value for *max*.

Displaying the number of files in a volume

To see how many files are in a volume and the maximum number of files allowed on the volume, complete the following step.

Step	Action
1	Enter the following command: **`maxfiles vol_name`** *vol_name* is the volume whose maximum number of files you are increasing. **Result:** A display like the following appears: `Volume home: maximum number of files is currently 120962 (2872 used)` **Note** The value returned reflects only the number of files that can be created by users; the inodes reserved for internal use are not included in this number.

Reallocating file and volume layout

About reallocation

If your volumes contain large files or LUNs that store information that is frequently accessed and revised (such as databases), the layout of your data can become suboptimal. Additionally, when you add disks to an aggregate, your data is no longer evenly distributed across all of the disks. The Data ONTAP `reallocate` commands allow you to reallocate the layout of files, LUNs or entire volumes for better data access.

For more information

For more information about the reallocation commands, see the *Block Access Management Guide for iSCSI* or the *Block Access Guide for FCP*, keeping in mind that for reallocation, files are managed exactly the same as LUNs.

Managing FlexCache volumes

About FlexCache volumes

A FlexCache volume is a sparsely populated volume on a local NetApp system backed by a volume on a different, possibly remote, NetApp system. A sparsely populated volume provides access to all data in the backing volume without requiring that the data be physically present in the sparsely populated volume.

You use FlexCache volumes to speed up access to remote data, or to offload traffic from heavily accessed volumes. Because the cached data must be ejected when the data is changed, FlexCache volumes work best for data that does not change often.

About this section

This section contains the following topics:

- "How FlexCache volumes work" on page 230
- "Sample FlexCache deployments" on page 236
- "Creating FlexCache volumes" on page 238
- "Sizing FlexCache volumes" on page 240
- "Administering FlexCache volumes" on page 242

How FlexCache volumes work

Direct access to cached data

A FlexCache volume is a sparsely populated volume on a local (caching) NetApp system that is backed by a volume on a different, possibly remote, (origin) NetApp system. A sparsely populated volume provides access to all data in the origin volume without requiring that the data be physically present in the sparsely populated volume.

When a client requests data from the FlexCache volume, the data is read through the network from the origin system and cached on the FlexCache volume. Subsequent requests for that data are then served directly from the FlexCache volume. In this way, clients in remote locations are provided with direct access to cached data. This improves performance when data is accessed repeatedly, because after the first request, the data no longer has to travel across the network.

FlexCache license requirement

You must have the `flex_cache` license installed before you can create FlexCache volumes. For more information about licensing, see the *System Administration Guide*.

Types of volumes you can use

A FlexCache volume must always be a FlexVol. FlexCache volumes can be created in the same aggregate as regular FlexVol volumes.

The origin volume can be a FlexVol or a traditional volume; it can also be a SnapLock volume. The origin volume cannot be a FlexCache volume itself, nor can it be a qtree.

Cache objects

The following objects can be cached in a FlexCache volume:

- files
- directories
- symbolic links

Note

In this document, the term “file” is used to refer to all of these object types.

File attributes are cached

When a data block from a specific file is requested from a FlexCache volume, then the attributes of that file are cached, and that file is considered to be cached. This is true even if not all of the data blocks that make up that file are present in the cache.

Cache consistency

Cache consistency for FlexCache volumes is achieved using three primary techniques: *delegations*, *attribute cache timeouts*, and *write operation proxy*.

Delegations: When data from a particular file is retrieved from the origin volume, the origin volume can give a delegation for that file to the caching volume. If that file is changed on the origin volume, whether from another caching volume or through direct client access, then the origin volume revokes the delegation for that file with all caching volumes that have that delegation. You can think of a delegation as a contract between the origin volume and the caching volume; as long as the caching volume has the delegation, the file has not changed.

Note

Delegations can cause a small performance decrease for writes to the origin volume, depending on the number of caching volumes holding delegations for the file being modified.

Delegations are not always used. The following list outlines situations when delegations cannot be used to guarantee that an object has not changed:

- Objects other than regular files do not use delegations

  Delegations are not used for any objects other than regular files. Directories, symbolic links, and other objects have no delegations.
- When connectivity is lost

  If connectivity is lost between the caching and origin filers, then delegations cannot be honored and must be considered to be revoked.
- When the maximum number of delegations has been reached

  If the origin volume cannot store all of its delegations, it might revoke an existing delegation to make room for a new one.

Attribute cache timeouts: When data is retrieved from the origin volume, the file that contains that data is considered valid in the FlexCache volume as long as a delegation exists for that file. However, if no delegation for the file exists, then it is considered valid for a specified length of time, called the attribute cache timeout. As long as a file is considered valid, if a client reads from that file and the requested data blocks are cached, the read request is fulfilled without any access to the origin volume.

If a client requests data from a file for which there are no delegations, and the attribute cache timeout has been exceeded, the FlexCache volume verifies that the attributes of the file have not changed on the origin system. Then one of the following actions is taken:

- If the attributes of the file have not changed since the file was cached, then the requested data is either directly returned to the client (if it was already in the FlexCache volume) or retrieved from the origin system and then returned to the client.
- If the attributes of the file have changed, the file is marked as invalid in the cache. Then the requested data blocks are read from the origin system, as if it were the first time that file had been accessed from that FlexCache volume.

With attribute cache timeouts, clients can get stale data when the following conditions are true:

- There are no delegations for the file on the caching volume
- The file's attribute cache timeout has not been reached
- The file has changed on the origin volume since it was last accessed by the caching volume

To prevent clients from ever getting stale data, you can set the attribute cache timeout to zero. However, this will negatively affect your caching performance, because then every data request for which there is no delegation causes an access to the origin system.

The attribute cache timeouts are determined using volume options. The volume option names and default values are outlined in the following table.

Volume option name	Description	Default Value
`acdirmax`	Attribute cache timeout for directories	30s
`acregmax`	Attribute cache timeout for regular files	30s
`acsymmax`	Attribute cache timeout for symbolic links	30s
`actimeo`	Attribute cache timeout for all objects	30s

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

Write operation proxy: If the client modifies the file, that operation is proxied through to the origin filer, and the file is ejected from the cache. This also changes the attributes of the file on the origin volume, so any other FlexCache volume that has that data cached will re-request the data once the attribute cache timeout is reached and a client requests that data.

Cache hits and misses

When a client makes a read request, if the relevant block is cached in the FlexCache volume, the data is read directly from the FlexCache volume. This is called a cache hit. Cache hits are the result of a previous request.

A cache hit can be one of the following types:

- Hit

  The requested data is cached and no verify is required; the request is fulfilled locally and no access to the origin filer is made.
- Hit-Verify

  The requested data is cached but the verification timeout has been exceeded, so the file attributes are verified against the origin system. No data is requested from the origin system.

If data is requested that is not currently on the FlexCache volume, or if that data has changed since it was cached, the caching system loads the data from the origin system and then returns it to the requesting client. This is called a cache miss.

A cache miss can be one of the following types:

- Miss

  The requested data is not in the cache; it is read from the origin system and cached.
- Miss-Verify

  The requested data is cached, but the file attributes have changed since the file was cached; the file is ejected from the cache and the requested data is read from the origin system and cached.

Limitations of FlexCache volumes

There are certain limitations of the FlexCache feature, for both the caching volume and for the origin volume.

Limitations of FlexCache caching volumes: You cannot use the following capabilities on FlexCache volumes (these limitations do *not* apply to the origin volumes):

- Client access using any protocol other than NFSv2 or NFSv3
- Snapshot creation
- SnapRestore
- SnapMirror (qtree or volume)
- SnapVault
- FlexClone creation
- `ndmp`
- Quotas
- Qtrees
- `vol copy`
- FlexCache volumes cannot be put into any vFiler other than vFiler0

Limitations of FlexCache origin volumes: You cannot perform the following operations on a FlexCache origin volume or NetApp system without rendering all FlexCache volumes backed by that origin volume unusable:

- You cannot move an origin volume between vFiler units or to vFiler0 using any of the following commands:
  - `vfiler move`
  - `vfiler add`
  - `vfiler remove`
  - `vfiler destroy`

  If you want to perform these operations on the origin volume, you can delete all FlexCache volumes backed by that volume, perform the operation, and then recreate the FlexCache volumes.

  Note

  You can use SnapMover (`vfiler migrate`) to migrate an origin volume without having to recreate any FlexCache volumes backed by that volume.

- You cannot change the IP address of the origin NetApp system.

  If you must change the IP address of the origin system, you can delete all FlexCache volumes backed by the volumes on that system, change the IP address, then recreate the FlexCache volumes.

What happens when connectivity to origin system is lost

If connectivity between the caching and origin NetApp systems is lost after a FlexCache volume is created, any data access that does not require access to the origin system succeeds. However, any operation that requires access to the origin volume, either because the requested data is not cached or because its attribute cache timeout has been exceeded, hangs until connectivity is reestablished.

Sample FlexCache deployments

WAN or LAN deployment

A FlexCache volume can be deployed in a WAN configuration or a LAN configuration.

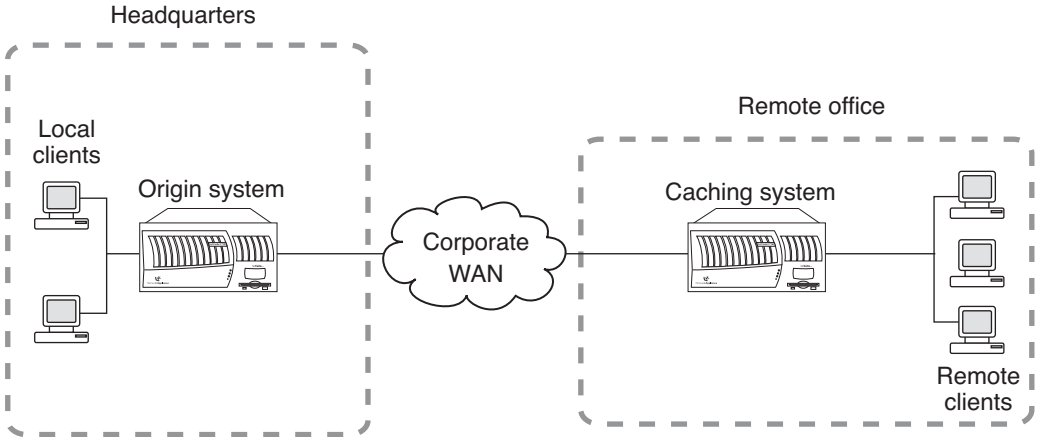
WAN deployment: In a WAN deployment, the FlexCache volume is remote from the data center. As clients request data, the FlexCache volume caches popular data, giving the end user faster access to information.

LAN deployment: In a LAN deployment, or accelerator mode, the FlexCache volume is local to the administrative data center, and is used to offload work from busy file servers and free system resources.

WAN deployment

In a WAN deployment, the FlexCache volume is placed as close as possible to the remote office. Client requests are then explicitly directed to the appliance. If valid data exists in the cache, that data is served directly to the client. If the data does not exist in the cache, it is retrieved across the WAN from the origin NetApp system, cached in the FlexCache volume, and returned to the client.

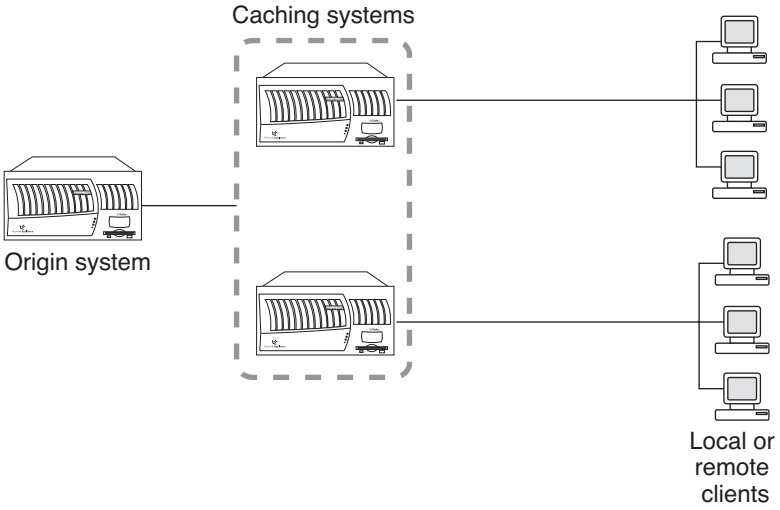
The following diagram shows a typical FlexCache WAN deployment.



LAN deployment

In a LAN deployment, a FlexCache volume is used to offload busy data servers. Frequently accessed data, or “hot objects” are replicated and cached by the FlexCache volume. This saves network bandwidth, reduces latency, and improves storage use, because only the most frequently used data is moved and stored.

The following example illustrates a typical LAN deployment.



Creating FlexCache volumes

Before creating a FlexCache volume

Before creating a FlexCache volume, ensure that you have the following configuration options set correctly:

- `flex_cache` license installed on the caching system
- `flexcache.access` option on origin system set to allow access from caching system

  For more information about this option, see the na_protocolaccess(8) man page.
- `flexcache.enable` option on the origin system set to `on`
- NFS licensed and enabled for the caching system
- Both the caching and origin systems running Data ONTAP 7.0.1 or later

Creating a FlexCache volume

To create a FlexCache volume, complete the following steps.

Step	Action
1	Ensure that your options are set correctly as outlined in “Before creating a FlexCache volume” on page 238.

Step	Action
2	Enter the following command: **`vol create cache_vol aggr size{k|m|g|t} -S origin:source_vol`** *cache_vol* is the name of the new FlexCache volume you want to create. *aggr* is the name of the containing aggregate for the new FlexCache volume. *size*{ `k` | `m` | `g` | `t` } specifies the FlexCache volume size in kilobytes, megabytes, gigabytes, or terabytes. For example, you would enter `20m` to indicate twenty megabytes. If you do not specify a unit, size is taken as bytes and rounded up to the nearest multiple of 4 KB. **Note** Because FlexCache volumes are sparsely populated, you can make the FlexCache volume smaller than the source volume. However, the larger the FlexCache volume is, the better caching performance it provides. For more information about sizing FlexCache volumes, see "Sizing FlexCache volumes" on page 240. *origin* is the name of the origin NetApp system *source_vol* is the name of the volume you want to use as the origin volume on the origin system. **Result:** The new FlexCache volume is created and an entry is added to the /etc/export file for the new volume. **Example:** The following command creates a 100-MB FlexCache volume called newcachevol, in the aggregate called aggr1, with a source volume vol1 on NetApp system corp_filer. `vol create newcachevol aggr1 100M -S corp_filer:vol1`

Sizing FlexCache volumes

About sizing FlexCache volumes

FlexCache volumes can be smaller than their origin volumes. However, making your FlexCache volume too small can negatively impact your caching performance. When the FlexCache volume begins to fill up, it flushes old data to make room for newly requested data. When that old data is requested again, it must be retrieved from the origin volume.

For best performance, create FlexCache volumes at least 135% of the size of your working set (data you expect to access frequently).

Note

Creating FlexCache volumes larger than 140% of the size of the source volume does not increase caching performance.

FlexCache volumes and space management

FlexCache volumes do not use space management in the same manner as regular FlexVol volumes. When you create a FlexCache volume of a certain size, that volume will not grow larger than that size. However, only a certain amount of space is preallocated for the volume. The amount of disk space allocated for a FlexCache volume is determined by the value of the `flexcache_min_reserved` volume option.

Note

The default value for the `flexcache_min_reserved volume` option is 100 MB. You should not need to change the value of this option.

Space allocation for multiple volumes in the same aggregate

You can have multiple FlexCache volumes in the same aggregate; you can also have regular FlexVol volumes in the same aggregate as your FlexCache volumes.

Multiple FlexCache volumes in the same aggregate: When you put multiple FlexCache volumes in the same aggregate, they can each be sized to be as large as the aggregate permits. This is because only the amount of space specified by the `flexcache_min_reserved` volume option is actually reserved for each one. The rest of the space is allocated as needed. This means that a “hot”

FlexCache volume, or one that is receiving more data accesses, is permitted to take up more space, while a FlexCache volume that is not being accessed as often will gradually be reduced in size.

FlexVol volumes and FlexCache volumes in the same aggregate: If you have regular FlexVol volumes in the same aggregate as your FlexCache volumes, and you start to fill up the aggregate, the FlexCache volumes can lose some of their unreserved space (only if they are not currently using it). In this case, when the FlexCache volume needs to fetch a new data block and it does not have enough free space to accommodate it, a data block must be ejected from one of the FlexCache volumes to make room for the new data block.

If this situation causes too many cache misses, you can add more space to your aggregate or move some of your data to another aggregate.

Using the df command with FlexCache volumes

When you use the `df` command on the caching NetApp system, you display the disk free space for the *origin* volume, rather than the local caching volume. You can display the disk free space for the local caching volume by using the `-L` option for the `df` command.

Administering FlexCache volumes

Viewing FlexCache statistics

Data ONTAP provides statistics about FlexCache volumes to help you understand the access patterns and administer the FlexCache volumes effectively. You can get statistics for your FlexCache volumes using the following commands:

- `flexcache stats` (client and server statistics)
- `nfsstat` (client statistics only)

For more information about these commands, see the na_flexcache(1) and nfsstat(1) man pages.

Client (caching system) statistics: You can use client statistics to see how how many operations are being served by the FlexCache rather than the origin system. A large number of cache misses after the FlexCache volume has had time to become populated may indicate that the FlexCache volume is too small and data is being discarded and fetched again later.

To view client FlexCache statistics, you use the `-C` option of the `flexcache stats` command on the *caching* system.

You can also view the `nfs` statistics for your FlexCache volumes using the `-C` option for the `nfsstat` command.

Server (origin system) statistics: You can use server statistics to see how much load is hitting the origin volume and which clients are causing that load. This can be useful if you are using the LAN deployment to offload an overloaded volume, and you want to make sure that the load is evenly distributed among the caching volumes.

To view server statistics, you use the `-S` option of the `flexcache stats` command on the *origin* system.

Note

You can also view the server statistics by client, using the `-c` option. The `flexcache.per_client_stats` option must be set to On.

Flushing files from FlexCache volumes

If you know that a specific file has changed on the origin volume and you want to flush it from your FlexCache volume before it is accessed, you can use the `flexcache eject` command. For more information about this command, see the na_flexcache(1) man page.

Space management for volumes and files

What space management is

The space management capabilities of Data ONTAP allow you to configure your NetApp systems to provide the storage availability required by the users and applications accessing the system, while using your available storage as effectively as possible.

Data ONTAP provides space management using the following capabilities:

- Space guarantees

  This capability is available only for FlexVol volumes.

  For more information, see "Space guarantees" on page 247.
- Space reservations

  For more information, see "Space reservations" on page 252 and the *Block Access Management Guide for iSCSI* or the *Block Access Management Guide for FCP*.
- Fractional reserve

  This capability is an extension of space reservations that is new for Data ONTAP 7.0.

  For more information, see "Fractional reserve" on page 254 and the *Block Access Management Guide for iSCSI* or the *Block Access Management Guide for FCP*.

Space management and files

Space reservations and fractional reserve are designed primarily for use with LUNs. Therefore, they are explained in greater detail in the *Block Access Management Guide for iSCSI* and the *Block Access Management Guide for FCP*. If you want to use these space management capabilities for files, consult those guides, keeping in mind that files are managed by Data ONTAP exactly the same as LUNs, except that space reservations are enabled for LUNs by default, whereas space reservations must be explicitly enabled for files.

What kind of space management to use

The following table can help you determine which space management capabilities best suit your requirements.

If…	Then use…	Typical usage	Notes
◆ You want management simplicity ◆ You have been using a version of Data ONTAP earlier than 7.0 and want to continue to manage your space the same way	◆ FlexVol volumes with space guarantee = `volume` ◆ Traditional volumes	NAS file systems	This is the easiest option to administer. As long as you have sufficient free space in the volume, writes to any file in this volume will always succeed. For more information about space guarantees, see "Space guarantees" on page 247.
◆ Writes to certain files must always succeed ◆ You want to overcommit your aggregate	◆ FlexVol volumes with space guarantee = `file` OR ◆ Traditional volume AND space reservation enabled for files that require writes to succeed	◆ LUNs ◆ Databases	This option enables you to guarantee writes to specific files. For more information about space guarantees, see "Space guarantees" on page 247. For more information about space reservations, see "Space reservations" on page 252 and the *Block Access Management Guide for iSCSI* or the *Block Access Management Guide for FCP*.

If…	Then use…	Typical usage	Notes
◆ You need even more effective storage usage than file space reservation provides ◆ You actively monitor available space on your volume and can take corrective action when needed ◆ Snapshots are short-lived ◆ Your rate of data overwrite is relatively predictable and low	◆ FlexVol volumes with space guarantee = `volume` OR ◆ Traditional volume AND Space reservation on for files that require writes to succeed AND Fractional reserve < 100%	◆ LUNs (with active space monitoring) ◆ Databases (with active space monitoring)	With fractional reserve <100%, it is possible to use up all available space, even with space reservations on. Before enabling this option, be sure either that you can accept failed writes or that you have correctly calculated and anticipated storage and snapshot usage. For more information, see “Fractional reserve” on page 254 and the *Block Access Management Guide for iSCSI* or the *Block Access Management Guide for FCP.*
◆ You want to overcommit your aggregate ◆ You actively monitor available space on your aggregate and can take corrective action when needed	FlexVol volumes with space guarantee = `none`	◆ Storage providers who need to provide storage that they know will not immediately be used ◆ Storage providers who need to allow available space to be dynamically shared between volumes	With an overcommitted aggregate, writes can fail due to insufficient space. For more information about aggregate overcommitment, see “Aggregate overcommitment” on page 250.

Space guarantees

What space guarantees are

Space guarantees on a FlexVol ensure that writes to a specified FlexVol or writes to files with space reservations enabled do not fail because of lack of available space in the containing aggregate.

Other operations such as creation of snapshots or new volumes in the containing aggregate can occur only if there is enough available uncommitted space in that aggregate; other operations are restricted from using space already committed to another volume.

When the uncommitted space in an aggregate is exhausted, only writes to volumes or files in that aggregate with space guarantees are guaranteed to succeed.

- A space guarantee of `volume` preallocates space in the aggregate for the volume. The preallocated space cannot be allocated to any other volume in that aggregate.

  The space management for a FlexVol with space guarantee of `volume` is equivalent to a traditional volume, or all volumes in versions of Data ONTAP earlier than 7.0.
- A space guarantee of `file` preallocates space in the volume so that any file in the volume with space reservation enabled can be completely rewritten, even if its blocks are pinned for a snapshot.

  For more information on file space reservation see "Space reservations" on page 252.
- A FlexVol with a space guarantee of `none` reserves no extra space; writes to LUNs or files contained by that volume could fail if the containing aggregate does not have enough available space to accommodate the write.

  Note

  Because out-of-space errors are unexpected in a CIFS environment, do not set space guarantee to `none` for volumes accessed using CIFS.

Space guarantee is an attribute of the volume. It is persistent across filer reboots, takeovers, and givebacks, but it does not persist through reversions to an earlier Data ONTAP software version.

Space guarantees are honored only for online volumes. If you take a volume offline, any committed but unused space for that volume becomes available for other volumes in that aggregate. When you bring that volume back online, if there is not sufficient available space in the aggregate to fulfill its space guarantees, you must use the force (`-f`) option, and the volume's space guarantees are disabled.

For more information, see "Bringing a volume online in an overcommitted aggregate" on page 251.

Traditional volumes and space management

Traditional volumes provide the same space guarantee as FlexVol volumes with space guarantee of `volume`. To guarantee that writes to a specific file in a traditional volume will always succeed, you need to enable space reservations for that file. (LUNs have space reservations enabled by default.)

For more information about space reservations, see "Space reservations" on page 252.

Specifying space guarantee at FlexVol creation time

To specify the space guarantee for a volume at creation time, complete the following steps.

Note

To create a FlexVol with space guarantee of `volume`, you can ignore the `guarantee` parameter, because `volume` is the default.

Step	Action
1	Enter the following command: **`vol create`** ***`f_vol_name aggr_name`*** **`-s {volume|file|none}`** ***`size`*****`{k|m|g|t}`** *f_vol_name* is the name for the new FlexVol (without the /vol/ prefix). This name must be unique from all other volume names on the filer. *aggr_name* is the containing aggregate for this FlexVol. `-s` specifies the space guarantee to be used for this volume. The possible values are `{volume|file|none}`. The default value is `volume`. *size* `{k|m|g|t}` specifies the maximum volume size in kilobytes, megabytes, gigabytes, or terabytes. For example, you would enter `4m` to indicate four megabytes. If you do not specify a unit, size is considered to be in bytes and rounded up to the nearest multiple of 4 KB.
2	To confirm that the space guarantee is set, enter the following command: **`vol options`** ***`f_vol_name`***

Changing space guarantee for existing volumes

To change the space guarantee for an existing FlexVol, complete the following steps.

Step	Action
1	Enter the following command: **`vol options f_vol_name guarantee guarantee_value`** *f_vol_name* is the name of the FlexVol whose space guarantee you want to change. *guarantee_value* is the space guarantee you want to assign to this volume. The possible values are `volume`, `file`, and `none`. **Note** If there is insufficient space in the aggregate to honor the space guarantee you want to change to, the command succeeds, but a warning message is printed and the space guarantee for that volume is disabled.
2	To confirm that the space guarantee is set, enter the following command: **`vol options f_vol_name`**

Aggregate overcommitment

Aggregate overcommitment provides flexibility to the storage provider. Using aggregate overcommitment, you can appear to provide more storage than is actually available from a given aggregate. This could be useful if you are asked to provide greater amounts of storage than you know will be used immediately. Alternatively, if you have several volumes that sometimes need to grow temporarily, the volumes can dynamically share the available space with each other.

To use aggregate overcommitment, you create FlexVol volumes with a space guarantee of `none` or `file`. With a space guarantee of `none` or `file`, the volume size is not limited by the aggregate size. In fact, each volume could, if required, be larger than the containing aggregate. The storage provided by the aggregate is used up only as LUNs are created or data is appended to files in the volumes.

Of course, when the aggregate is overcommitted, it is possible for these types of writes to fail due to lack of available space:

- Writes to any volume with space guarantee of `none`
- Writes to any file that does not have space reservations enabled and that is in a volume with space guarantee of `file`

Therefore, if you have overcommitted your aggregate, you must monitor your available space and add storage to the aggregate as needed to avoid write errors due to insufficient space.

Note

Because out-of-space errors are unexpected in a CIFS environment, do not set space guarantee to `none` for volumes accessed using CIFS.

Bringing a volume online in an overcommitted aggregate

When you take a FlexVol offline, it relinquishes its allocation of storage space in its containing aggregate. Storage allocation for other volumes in that aggregate while that volume is offline can result in that storage being used. When you bring the volume back online, if there is insufficient space in the aggregate to fulfill the space guarantee of that volume, the normal `online` command fails unless you force the volume online by using the `-f` flag.

Caution

When you force a FlexVol to come online due to insufficient space, the space guarantees for that volume are disabled. That means that attempts to write to that volume could fail due to insufficient available space. In environments that are sensitive to that error, such as CIFS or LUNs, forcing a volume online should be avoided if possible.

When you make sufficient space available to the aggregate, the space guarantees for the volume are automatically re-enabled.

To bring a FlexVol online when there is insufficient storage space to fulfill its space guarantees, complete the following step.

Step	Action
1	Enter the following command: **`vol online vol_name -f`** *vol_name* is the name of the volume you want to force online.

Space reservations

What space reservations are

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

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

File space reservation is an attribute of the file; it is persistent across filer reboots, takeovers, and givebacks.

There is no way to automatically enable space reservations for every file in a given volume, as you could with versions of Data ONTAP earlier than 7.0 using the `create_reserved` option. In Data ONTAP 7.0, to guarantee that writes to a specific file will always succeed, you need to enable space reservations for that file. (LUNs have space reservations enabled by default.)

Note

For more information about using space reservation for LUNs, see the *Block Access Management Guide for FCP* or the *Block Access Management Guide for iSCSI*.

For more information about using space reservation for files, see the *Block Access Management Guide for FCP* or the *Block Access Management Guide for iSCSI*, keeping in mind that Data ONTAP manages files exactly the same as LUNs, except that space reservations are enabled automatically for LUNs, whereas for files, you must explicitly enable space reservations.

Enabling space reservation for a specific file

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

Step	Action
1	Enter the following command: **`file reservation file_name [enable|disable]`** *file_name* is the file in which file space reservation is set. `enable` turns space reservation on for the file *file_name.* `disable` turns space reservation off for the file *file_name*. **Example:** `file reservation myfile enable` **Note** In FlexVol volumes, the volume option `guarantee` must be set to `file` or `volume` for file space reservations to work. For more information, see "Space guarantees" on page 247.

Turning on space reservation for a file fails if there is not enough available space for the new reservation.

Querying space reservation for files

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

Step	Action
1	Enter the following command: **`file reservation file_name`** *file_name* is the file you want to query the space reservation status for. **Example:** `file reservation myfile` **Result:** The space reservation status for the specified file is displayed: `space reservations for file /vol/flex1/1gfile: off`

Fractional reserve

Fractional reserve

If you have enabled space reservation for a file or files, you can reduce the space that you preallocate for those reservations using fractional reserve. Fractional reserve is an option on the volume, and it can be used with either traditional or FlexVol volumes. Setting fractional reserve to less than 100 causes the space reservation held for all space-reserved files in that volume to be reduced to that percentage. Writes to the space-reserved files are no longer unequivocally guaranteed; you must monitor your reserved space and take action if your free space becomes scarce.

Fractional reserve is generally used for volumes that hold LUNs with a small percentage of data overwrite.

Note

If you are using fractional reserve in environments where write errors due to lack of available space are unexpected, you must monitor your free space and take corrective action to avoid write errors.

For more information about fractional reserve, see the *Block Access Management Guide for iSCSI* or the *Block Access Management Guide for FCP.*

Qtree Management 7

About this chapter

This chapter describes how to use qtrees to manage user data. Read this chapter if you plan to organize user data into smaller units (qtrees) for flexibility or in order to use tree quotas.

Topics in this chapter

This chapter discusses the following topics:

- "Understanding qtrees" on page 256
- "Understanding qtree creation" on page 258
- "Creating qtrees" on page 260
- "Understanding security styles" on page 261
- "Changing security styles" on page 263
- "Changing the CIFS oplocks setting" on page 265
- "Displaying qtree status" on page 268
- "Displaying qtree access statistics" on page 269
- "Converting a directory to a qtree" on page 270
- "Renaming or deleting qtrees" on page 273

Additional qtree operations are described in other chapters or other guides:

- For information about setting usage quotas for users, groups, or qtrees, see the chapter titled "Quota Management" on page 275.
- For information about configuring and managing qtree-based SnapMirror replication, see the *Data Protection Online Backup and Recovery Guide*.

Understanding qtrees

What qtrees are

A qtree is a logically defined file system that can exist as a special subdirectory of the root directory within either a traditional volume or a FlexVol.

Note

You can have a maximum of 4,995 qtrees on any volume.

When creating qtrees is appropriate

You might create a qtree for either or both of the following reasons:

- You can easily create qtrees for managing and partitioning your data within the volume.
- You can create a qtree to assign user- or workgroup-based soft or hard usage quotas to limit the amount of storage space that a specified user or group of users can consume on the qtree to which they have access.

Qtrees and volumes comparison

In general, qtrees are similar to volumes. However, they have the following key differences:

- Snapshots can be enabled or disabled for individual volumes, but not for individual qtrees.
- Qtrees do not support space reservations or space guarantees.

Qtrees, traditional volumes, and FlexVol volumes have other differences and similarities as shown in the following table.

Function	Traditional volume	FlexVol	Qtree
Enables organizing user data	Yes	Yes	Yes
Enables grouping users with similar needs	Yes	Yes	Yes
Can assign a security style to determine whether files use UNIX or Windows NT permissions.	Yes	Yes	Yes

Function	Traditional volume	FlexVol	Qtree
Can configure the oplocks setting to determine whether files and directories use CIFS opportunistic locks.	Yes	Yes	Yes
Can be used as units of SnapMirror backup and restore operations	Yes	Yes	Yes
Can be used as units of SnapVault backup and restore operations	No	No	Yes
Easily expandable and shrinkable	No (expandable but not shrinkable)	Yes	Yes
Snapshots	Yes	Yes	No (qtree replication extractable from volume snapshots)
Manage user based quotas	Yes	Yes	Yes
Cloneable	No	Yes	No (but can be part of a FlexClone)

Understanding qtree creation

Qtree grouping criteria

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

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

Using qtrees for projects

One way to group files is to set up a qtree for a project, such as one maintaining a database. Setting up a qtree for a project provides you with the following capabilities:

- Set the security style of the project without affecting the security style of other projects.

  For example, you use NTFS-style security if the members of the project use Windows files and applications. Another project in another qtree can use UNIX files and applications, and a third project can use Windows as well as UNIX files.
- If the project uses Windows, set CIFS oplocks (opportunistic locks) as appropriate to the project, without affecting other projects.

  For example, if one project uses a database that requires no CIFS oplocks, you can set CIFS oplocks to Off on that project qtree. If another project uses CIFS oplocks, it can be in another qtree that has oplocks set to On.
- Use quotas to limit the disk space and number of files available to a project qtree so that the project does not use up resources that other projects and users need. For instructions about managing disk space by using quotas, see Chapter 8, “Quota Management,” on page 275.
- Back up and restore all the project files as a unit.

Using qtrees for backups

You can back up individual qtrees to

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

Detailed information

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

- "Creating qtrees" on page 260
- "Understanding security styles" on page 261

If you do not want to accept the default security style of a volume or a qtree, you can change it, as described in "Changing security styles" on page 263.

If you do not want to accept the default CIFS oplocks setting of a volume or a qtree, you can change it, as described in "Changing the CIFS oplocks setting" on page 265.

Creating qtrees

Creating a qtree

To create a qtree, complete the following step.

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

Examples:

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

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

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

```
qtree create news
```

Understanding security styles

About security styles

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

Note

Although security styles can be applied to both qtrees and volumes, they are not shown as a volume attribute, and are managed for both volumes and qtrees using the `qtree` command.

Security styles

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

Security style	Description	Effect of changing to the style
NTFS	Exactly like Windows NT NTFS: Files and directories have Windows NT file-level permission settings. **Note** To use NTFS security, the filer must be licensed for CIFS.	If the change is from a mixed qtree, Windows NT permissions determine file access for a file that had Windows NT permissions. Otherwise, UNIX-style permission bits determine file access for files created before the change. **Note** If the change is from a CIFS filer to a multiprotocol filer, as described in "Parameters to accept or change after volume creation" on page 185, and the /etc directory is a qtree, its security style changes to NTFS.

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

Note

When you create an NTFS qtree or change a qtree to NTFS, every Windows user is given full access to the qtree, by default. You must change the permissions if you want to restrict access to the qtree for some users. If you do not set NTFS file security on a file, UNIX permissions are enforced.

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

Changing security styles

When to change the security style of a qtree or volume

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

- You might want to change the security style of a qtree after creating it to match the needs of the users of the qtree.
- You might want to change the security style to accommodate other users or files. For example, if you start with an NTFS qtree and subsequently want to include UNIX files and users, you might want to change the qtree from an NTFS qtree to a mixed qtree.

Effects of changing the security style on quotas

Changing the security style of a qtree or volume requires quota reinitialization if quotas are in effect. For information about how changing the security style affects quota calculation, see "Turning quota message logging on or off" on page 314.

Changing the security style of a qtree

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

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

Step	Action
2	If you have quotas in effect on the qtree whose security style you just changed, reinitialize quotas on the volume containing this qtree. **Result:** This allows Data ONTAP to recalculate the quota usage for users who own files with ACL or UNIX security on this qtree. For information about reinitializing quotas, see “Activating or reinitializing quotas” on page 306.

Caution

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

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

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

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

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

Changing the CIFS oplocks setting

What CIFS oplocks do

CIFS oplocks (opportunistic locks) enable the redirector on a CIFS client in certain file-sharing scenarios to perform client-side caching of read-ahead, write-behind, and lock information. A client can then work with a file (read or write it) without regularly reminding the server that it needs access to the file in question. This improves performance by reducing network traffic.

For more information on CIFS oplocks, see the CIFS section of the *File Access and Protocols Management Guide*.

When to turn CIFS oplocks off

CIFS oplocks on the filer are on by default.

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

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

Otherwise, you can leave CIFS oplocks on.

Effect of the cifs.oplocks.enable option

The `cifs.oplocks.enable` option enables and disables CIFS oplocks for the entire filer.

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

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

Enabling CIFS oplocks for a specific volume or qtree

To enable CIFS opslocks for a specific volume or a qtree, complete the following steps.

Step	Action
1	Make sure the global `cifs.oplocks.enable` option is set to On.
2	Enter the following command: **`qtree oplocks`** ***`path`*** **`enable`** *path* is the path name of the volume or the qtree.
3	To verify that CIFS oplocks were updated as expected, enter the following command: **`qtree status`** ***`vol_name`*** *vol_name* is the name of the specified volume, or the volume that contains the specified qtree.

Example: To enable CIFS oplocks on the proj1 qtree in vol2, use the following commands:

```
filer1> options cifs.oplocks.enable on
filer1> qtree oplocks /vol/vol2/proj enable
```

Disabling CIFS oplocks for a specific volume or qtree

To disable CIFS opslocks for a specific volume or a qtree, complete the following steps.

Caution

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

Step	Action
1	Enter the following command: **`qtree oplocks`** ***`path`*** **`disable`** *path* is the path name of the volume or the qtree.

Step	Action
2	To verify that CIFS oplocks were updated as expected, enter the following command: **`qtree status vol_name`** *vol_name* is the name of the specified volume, or the volume that contains the specified qtree.

Example: To disable CIFS oplocks on the proj1 qtree in vol2, use the following command:

```
qtree oplocks /vol/vol2/proj disable
```

Displaying qtree status

Determining the status of qtrees

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

Step	Action
1	Enter the following command: **`qtree status [-i] [-v] [`*`path`*`]`** The `-i` option includes the qtree ID number in the display. The `-v` option includes the owning vfiler, if the MultiStore license is enabled.

Example 1:

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

Example 2:

```
toaster> qtree status -v vol1
Volume    Tree    Style  Oplocks  Status        Owning vfiler
--------  -----   -----  -------- ------        -------------
vol1              unix   enabled  normal        vfiler0
vol1      engr    ntfs   disabled normal        vfiler0
vol1      backup  unix   enabled  snapmirrored  vfiler0
```

Example 3:

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

Displaying qtree access statistics

About qtree stats

The `qtree stats` command enables you to display statistics on user accesses to files in qtrees on your system. This can help you determine what qtrees are incurring the most traffic. Determining traffic patterns helps with qtree-based load balancing.

How the qtree stats command works

The `qtree stats` command displays the number of NFS and CIFS accesses to the designated qtrees since the counters were last reset. The qtree stats counters are reset when one of the following actions occurs:

- The system is booted.
- The volume containing the qtree is brought online.
- The counters are explicitly reset using the `qtree stats -z` command.

Using qtree stats

To use the `qtree stats` command, complete the following step.

Step	Action
1	Enter the following command: **`qtree stats [ -z ] [ path ]`** The `-z` option clears the counter for the designated qtree, or clears all counters if no qtree is specified.

Example:

```
toaster> qtree stats vol1
Volume        Tree         NFS ops       CIFS ops
--------      --------     -------       --------
vol1          proj1           1232             23
vol1          proj2             55            312
```

Example with -z option:

```
toaster> qtree stats -z vol1
Volume        Tree         NFS ops       CIFS ops
--------      --------     -------       --------
vol1          proj1              0              0
vol1          proj2              0              0
```

Converting a directory to a qtree

Converting a rooted directory to a qtree

A rooted directory is a directory at the root of a volume. If you have a rooted directory that you want to convert to a qtree, you must migrate the data contained in the directory to a new qtree with the same name, using your client application. The following process outlines the tasks you need to complete to convert a rooted directory to a qtree:

Stage	Task
1	Rename the directory to be made into a qtree.
2	Create a new qtree with the original directory name.
3	Use the client application to move the contents of the directory into the new qtree.
4	Delete the now-empty directory.

Note

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

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

Note

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

Converting a rooted directory to a qtree using a Windows client

To convert a rooted directory to a qtree using a Windows client, complete the following steps.

Step	Action
1	Open Windows Explorer.
2	Click the folder representation of the directory you want to change.

Step	Action
3	From the File menu, select Rename to give this directory a different name.
4	On the filer, use the `qtree create` command to create a new qtree with the original name.
5	In Windows Explorer, open the renamed folder and select the files inside it.
6	Drag these files into the folder representation of the new qtree. **Note** The more subfolders contained in a folder that you are moving across qtrees, the longer the move operation for that folder will take.
7	From the File menu, select Delete to delete the renamed, now-empty directory folder.

Converting a rooted directory to a qtree using a UNIX client

To convert a rooted directory to a qtree using a UNIX client, complete the following steps.

Step	Action
1	Open a UNIX window.
2	Use the `mv` command to rename the directory. **Example:** `client: mv /n/joel/vol1/dir1 /n/joel/vol1/olddir`
3	From the filer, use the `qtree create` command to create a qtree with the original name. **Example:** `filer: qtree create /n/joel/vol1/dir1`

Step	Action
4	From the client, use the `mv` command to move the contents of the old directory into the qtree. **Example:** `client: mv /n/joel/vol1/olddir/* /n/joel/vol1/dir1` **Note** The more subdirectories contained in a directory that you are moving across qtrees, the longer move operation for that directory will take.
5	Use the `rmdir` command to delete the old, now-empty directory. **Example:** `client: rmdir /n/joel/vol1/olddir`

Renaming or deleting qtrees

Before renaming or deleting a qtree

Before you rename or delete a qtree, ensure that the following conditions are true:

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

Renaming a qtree

To rename a qtree, complete the following steps.

Step	Action
1	Find the qtree you want to rename. **Note** The qtree appears as a normal directory at the root of the volume.
2	Rename the qtree using the method appropriate for your client. **Example:** The following command on a UNIX host renames a qtree: `mv old_name new_name` **Note** On a Windows host, rename a qtree by using Windows Explorer. If you have quotas on the renamed qtree, update the /etc/quotas file to use the new qtree name.

Deleting a qtree

To delete a qtree, complete the following steps.

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

Quota Management 8

About this chapter

This chapter describes how to restrict and track the disk space and number of files used by a user, group, or qtree.

Topics in this chapter

This chapter discusses the following topics:

- "Understanding quotas" on page 276
- "When quotas take effect" on page 279
- "Understanding default quotas" on page 280
- "Understanding derived quotas" on page 281
- "How Data ONTAP identifies users for quotas" on page 284
- "Notification when quotas are exceeded" on page 287
- "Understanding the /etc/quotas file" on page 288
- "Activating or reinitializing quotas" on page 306
- "Modifying quotas" on page 309
- "Deleting quotas" on page 312
- "Turning quota message logging on or off" on page 314
- "Effects of qtree changes on quotas" on page 316
- "Understanding quota reports" on page 318

For information about quotas and their effect in a client environment, see the *File Access and Protocols Management Guide*.

Understanding quotas

Reasons for specifying quotas

You specify a quota for the following reasons:

- To limit the amount of disk space or the number of files that can be used by a quota target
- To track the amount of disk space or the number of files used by a quota target, without imposing a limit
- To warn users when their disk space or file usage is high

Quota targets

A quota target can be

- A user, as represented by a UNIX ID or a Windows ID.
- A group, as represented by a UNIX group name or GID. Data ONTAP does not apply group quotas based on Windows IDs.
- A qtree, as represented by the path name to the qtree.

The quota target determines the quota type, as shown in the following table.

Quota target	Quota type
user	user quota
group	group quota
qtree	tree quota

Tree quotas

If you apply a tree quota to a qtree, the qtree is similar to a disk partition, except that you can change its size at any time. When applying a tree quota, Data ONTAP limits the disk space and number of files regardless of the owner of the disk space or files in the qtree. No users, including root and members of the BUILTIN\Administrators group, can write to the qtree if the write causes the tree quota to be exceeded.

Quota specifications

Quota specifications are stored in the /etc/quotas file, which you can edit at any time.

User and group quotas are applied on a per-volume or per-qtree basis. You cannot specify a single quota for an aggregate or for multiple volumes.

Example: You can specify that a user named jsmith can use up to 10 GB of disk space in the cad volume, or that a group named engineering can use up to 50 GB of disk space in the /vol/cad/projects qtree.

Explicit quotas

If the quota specification references the name or ID of the quota target, the quota is an explicit quota. For example, if you specify a user name, jsmith, as the quota target, the quota is an explicit user quota. If you specify the path name of a qtree, /vol/cad/engineering, as the quota target, the quota is an explicit tree quota.

For examples of explicit quotas, see "Explicit quota examples" on page 298.

Default quotas and derived quotas

The disk space used by a quota target can be restricted or tracked even if you do not specify an explicit quota for it in the /etc/quotas file. If a quota is applied to a target and the name or ID of the target does not appear in an /etc/quotas entry, the quota is called a derived quota.

For more information about default quotas, see "Understanding default quotas" on page 280. For more information about derived quotas, see "Understanding derived quotas" on page 281. For examples, see "Default quota examples" on page 299.

Hard quotas, soft quotas, and threshold quotas

A hard quota is a limit that cannot be exceeded. If an operation, such as a write, causes a quota target to exceed a hard quota, the operation fails. When this happens, a warning message is logged to the filer console and an SNMP trap is issued.

A soft quota is a limit that can be exceeded. When a soft quota is exceeded, a warning message is logged to the filer console and an SNMP trap is issued. When the soft quota limit is no longer being exceeded, another syslog message and SNMP trap are generated. You can specify both hard and soft quota limits for the amount of disk space used and the number of files created.

A threshold quota is similar to a soft quota. When a threshold quota is exceeded, a warning message is logged to the filer console and an SNMP trap is issued.

A single type of SNMP trap is generated for all types of quota events. You can find details on SNMP traps in the filer's /etc/mib/netapp.mib file.

Syslog messages about quotas contain qtree ID numbers rather than qtree names. You can correlate qtree names to the qtree ID numbers in syslog messages by using the `qtree status -i` command.

Tracking quotas

You can use tracking quotas to track, but not limit, the resources used by a particular user, group, or qtree. To see the resources used by that user, group, or qtree, you can use quota reports.

For examples of tracking quotas, see "Tracking quota examples" on page 298.

When quotas take effect

Prerequisite for quotas to take effect

You must activate quotas on a per-volume basis before Data ONTAP applies quotas to quota targets. For more information about activating quotas, see "Activating or reinitializing quotas" on page 306.

Note

Quota activation persists across halts and reboots. You should not activate quotas in the /etc/rc file.

About quota initialization

After you turn on quotas, Data ONTAP performs quota initialization. This involves scanning the entire file system in a volume and reading from the /etc/quotas file to compute the disk usage for each quota target.

Quota initialization is necessary under the following circumstances:

- You add an entry to the /etc/quotas file, but the quota target for that entry is not currently tracked by the filer.
- You change user mapping in the /etc/usermap.cfg file and you use the QUOTA_PERFORM_USER_MAPPING entry in the /etc/quotas file. For more information about QUOTA_PERFORM_USER_MAPPING, see "Special entries for mapping users" on page 301.
- You change the security style of a qtree from UNIX to either mixed or NTFS.
- You change the security style of a qtree from mixed or NTFS to UNIX.

Quota initialization can take a few minutes. The amount of time required depends on the size of the file system. During quota initialization, data access is not affected. However, quotas are not enforced until initialization completes.

For more information about quota initialization, see "Activating or reinitializing quotas" on page 306.

About changing a quota size

You can change the size of a quota that is being enforced. Resizing an existing quota, whether it is an explicit quota specified in the /etc/quotas file or a derived quota, does not require quota initialization. For more information about changing the size of a quota, see "Modifying quotas" on page 309.

Understanding default quotas

About default quotas

You can create a default quota for users, groups, or qtrees. A default quota applies to quota targets that are not explicitly referenced in the /etc/quotas file. You create default quotas by using an asterisk (*) in the Quota Target field in the /etc/quota file. For more information about creating default quotas, see “Fields of the /etc/quotas file” on page 292 and “Tracking quota examples” on page 298.

How to override a default quota

If you do not want Data ONTAP to apply a default quota to a particular target, you can create an entry in the /etc/quotas file for that target. The explicit quota for that target overrides the default quota.

Where default quotas are applied

You apply a default user or group quota on a per-volume or per-qtree basis.

You apply a default tree quota on a per-volume basis. For example, you can specify that a default tree quota be applied to the cad volume, which means that all qtrees created in the cad volume are subject to this quota but that qtrees in other volumes are unaffected.

Typical default quota usage

As an example, suppose you want a user quota to be applied to most users of your system. Rather than applying that quota individually to every user, you can create a default user quota that will be automatically applied to every user. If you want to change that quota for a particular user, you can override the default quota for that user by creating an entry for that user in the /etc/quotas file.

For an example of a default quota, see “Tracking quota examples” on page 298.

About default tracking quotas

If you do not want to specify a default user, group or tree quota limit, you can specify default tracking quotas. These special default quotas do not enforce any resource limits, but they enable you to resize rather than reinitialize quotas after adding or deleting quota file entries.

Understanding derived quotas

About derived quotas

Data ONTAP derives the quota information from the default quota entry in the /etc/quotas file and applies it if a write request affects the disk space or number of files used by the quota target. A quota applied due to a default quota, not due to an explicit entry in the /etc/quotas file, is referred to as a derived quota.

Derived user quotas from a default user quota

When a default user quota is in effect, Data ONTAP applies derived quotas to all users in the volume or qtree to which the default quota applies, except those users who have explicit entries in the /etc/quotas file. Data ONTAP also tracks disk usage for the root user and BUILTIN\Administrators in that volume or qtree.

Example: A default user quota entry specifies that users in the cad volume are limited to 10 GB of disk space and a user named jsmith creates a file in that volume. Data ONTAP applies a derived quota to jsmith to limit that user's disk usage in the cad volume to 10 GB.

Derived group quotas from a default group quota

When a default group quota is in effect, Data ONTAP applies derived quotas for all UNIX groups in the volume or qtree to which the quota applies, except those groups that have explicit entries in the /etc/quotas file. Data ONTAP also tracks disk usage for the group with GID 0 in that volume or qtree.

Example: A default group quota entry specifies that groups in the cad volume are limited to 10 GB of disk space and a file is created that is owned by a group named writers. Data ONTAP applies a derived quota to the writers group to limit its disk usage in the cad volume to 10 GB.

Derived tree quotas from a default tree quota

When a default tree quota is in effect, derived quotas apply to all qtrees in the volume to which the quota applies, except those qtrees that have explicit entries in the /etc/quotas file.

Example: A default tree quota entry specifies that qtrees in the cad volume are limited to 10 GB of disk space and a qtree named projects is created in the cad volume. Data ONTAP applies a derived quota to the cad projects qtree to limit its disk usage to 10 GB.

Default user or group quotas derived from default tree quotas

When a qtree is created in a volume that has a default tree quota defined in the /etc/quotas file, and that default quota is applied as a derived quota to the qtree just created, Data ONTAP also applies derived default user and group quotas to that qtree.

- If a default user quota or group quota is already defined for the volume containing the newly created qtree, Data ONTAP automatically applies that quota as the derived default user quota or group quota for that qtree.
- If no default user quota or group quota is defined for the volume containing the newly created qtree, then the effective derived user or group quota for that qtree is unlimited. In theory, a single user with no explicit user quota defined can use up the newly defined qtree's entire qtree quota allotment.
- You can replace the initial derived default user quotas or group quotas that Data ONTAP applies to the newly created qtree. To do so, you add explicit or default user or group quotas for the qtree just created to the /etc/quotas file.

Example of a default user quota for a volume applied to a qtree:

Suppose the default user quota in the cad volume specifies that each user is limited to 10 GB of disk space, and the default tree quota in the cad volume specifies that each qtree is limited to 100 GB of disk space. If you create a qtree named projects in the cad volume, a default tree quota limits the projects qtree to 100 GB. Data ONTAP also applies a derived default user quota, which limits to 10 GB the amount of space used by each user who does not have an explicit user quota defined in the /vol/cad/projects qtree.

You can change the limits on the default user quota for the /vol/cad/projects qtree or add an explicit quota for a user in the /vol/cad/projects qtree by using the `quota resize` command.

Example of no default user quota for a volume applied to a qtree:

If no default user quota is defined for the cad volume, and the default tree quota for the cad volume specifies that all qtrees are limited to 100 GB of disk space, and if you create a qtree named projects, Data ONTAP does not apply a derived default user quota that limits the amount of disk space that users can use on the /vol/cad/projects tree quota. In theory, a single user with no explicit user quota defined can use all 100 GB of a qtree's quota if no other user writes to disk space on the new qtree first.

In addition, UID 0, BUILTIN\Administrators, and GID 0 have derived quotas. These derived quotas do not limit the disk space and the number of files. They only track the disk space and the number of files owned by these IDs.

Even with no default user quota defined, no user with files on a qtree can use more disk space in that qtree than is allotted to that qtree as a whole.

Advantages of specifying default quotas

Specifying default quotas offers the following advantages:

- You can automatically apply a limit to a large set of quota targets without typing multiple entries in the /etc/quotas file. For example, if you want no user to use more than 10 GB of disk space, you can specify a default user quota of 10 GB of disk space instead of creating an entry in the /etc/quotas file for each user.
- You can be flexible in changing quota specifications. Because Data ONTAP already tracks disk and file usage for quota targets of derived quotas, you can change the specifications of these derived quotas without having to perform a full quota reinitialization.

  For example, you can create a default user quota for the vol1 volume that limits each user to 10 GB of disk space, and default tracking group and tree quotas for the cad volume. After quota initialization, these default quotas and their derived quotas go into effect.

  If you later decide that a user named jsmith should have a larger quota, you can add an /etc/quotas entry that limits jsmith to 20 GB of disk space, overriding the default 10-GB limit. After making the change to the /etc/quotas file, to make the jsmith entry effective, you can simply resize the quota, which takes less time than quota reinitialization.

  Without the default user, group and tree quotas, the newly created jsmith entry requires a full quota reinitialization to be effective.

How Data ONTAP identifies users for quotas

Two types of user IDs

When applying a user quota, Data ONTAP distinguishes one user from another based on the ID, which can be a UNIX ID or a Windows ID.

Format of a UNIX ID

If you want to apply user quotas to UNIX users, specify the UNIX ID of each user in one of the following formats:

- The user name, as defined in the /etc/passwd file or the NIS password map, such as jsmith.
- The UID, such as 20.
- A file or directory whose UID matches the user. In this case, you should choose a path name that will last as long as the user account remains on the system.

Note

Specifying a file or directory name only enables Data ONTAP to obtain the UID. Data ONTAP does not apply quotas to the file or directory, or to the volume in which the file or directory resides.

Restrictions on UNIX user names: A UNIX user name must not include a backslash (\) or an @ sign, because Data ONTAP treats names containing these characters as Windows names.

Special UID: You cannot impose restrictions on a user whose UID is 0. You can specify a quota only to track the disk space and number of files used by this UID.

Format of a Windows ID

If you want to apply user quotas to Windows users, specify the Windows ID of each user in one of the following formats:

- A Windows name specified in pre-Windows 2000 format. For details, see the section on specifying a Windows name in the CIFS chapter of the *File Access and Protocols Management Guide*.

  If the domain name or user name contains spaces or special characters, the entire Windows name must be in quotation marks, such as “tech support\john#smith”.
- A security ID (SID), as displayed by Windows in text form, such as S-1-5-32-544.

- A file or directory that has an ACL owned by the SID of the user. In this case, you should choose a path name that will last as long as the user account remains on the system.

  Note

  For Data ONTAP to obtain the SID from the ACL, the ACL must be valid.

  If a file or directory exists in a UNIX-style qtree or if the filer uses UNIX mode for user authentication, Data ONTAP applies the user quota to the user whose UID matches that of the file or directory, not to the SID.

How Windows group IDs are treated

Data ONTAP does not support group quotas based on Windows group IDs. If you specify a Windows group ID as the quota target, the quota is treated like a user quota.

The following list describes what happens if the quota target is a special Windows group ID:

- If the quota target is the Everyone group, a file whose ACL shows that the owner is Everyone is counted under the SID for Everyone.
- If the quota target is BUILTIN\Administrators, the entry is considered a user quota for tracking only. You cannot impose restrictions on BUILTIN\Administrators. If a member of BUILTIN\Administrators creates a file, the file is owned by BUILTIN\Administrators and is counted under the SID for BUILTIN\Administrators.

How quotas are applied to users with multiple IDs

A user can be represented by multiple IDs. You can set up a single user quota entry for such a user by specifying a list of IDs as the quota target. A file owned by any of these IDs is subject to the restriction of the user quota.

Example: A user has the UNIX UID 20 and the Windows IDs corp\john_smith and engineering\jsmith. For this user, you can specify a quota where the quota target is a list of the UID and Windows IDs. When this user writes to the filer, the specified quota applies, regardless of whether the write originates from UID 20, corp\john_smith, or engineering\jsmith.

Note

Quota targets listed in different quota entries are considered separate targets, even though the IDs belong to the same user.

Example: You can specify one quota that limits UID 20 to 1 GB of disk space and another quota that limits corp\john_smith to 2 GB of disk space, even though both IDs represent the same user. Data ONTAP applies quotas to UID 20 and corp\john_smith separately.

If the user has another Windows ID, engineering\jsmith, and there is no applicable quota entry (including a default quota), files owned by engineering\jsmith are not subject to restrictions, even though quota entries are in effect for UID 20 and corp\john_smith.

Root users and quotas

A root user is subject to tree quotas, but not user quotas or group quotas.

When root carries out a file or directory ownership change or other operation (such as the UNIX `chown` command) on behalf of a nonroot user, Data ONTAP checks the quotas based on the new owner but does not report errors or stop the operation even if the nonroot user's hard quota restrictions are exceeded. The root user can therefore carry out operations for a nonroot user (such as recovering data), even if those operations temporarily result in that nonroot user's quotas being exceeded.

Once the ownership transfer is carried out, however, a client system will report a disk space error for the nonroot user who is attempting to allocate more disk space while the quota is still exceeded.

Notification when quotas are exceeded

Console messages

When Data ONTAP receives a write request, it first determines whether the file to be written is in a qtree. If it is, and the write would exceed any hard quota, the write fails and a message is written to the console describing the type of quota exceeded and the volume. If the write would exceed any soft quota, the write succeeds, but a message is still written to the console.

SNMP notification

SNMP traps can be used to arrange e-mail notification when hard or soft quotas are exceeded. You can access and adapt a sample quota notification script on the NOW site at http://now.netapp.com/ under Software Downloads, in the Tools and Utilities section.

Understanding the /etc/quotas file

About this section

This section provides information about the /etc/quotas file so that you can specify user, group, or tree quotas.

Detailed information

This section discusses the following topics:

- "Overview of the /etc/quotas file" on page 289
- "Fields of the /etc/quotas file" on page 292
- "Sample quota entries" on page 298
- "Special entries for mapping users" on page 301
- "How disk space owned by default users is counted" on page 305

Overview of the /etc/quotas file

Contents of the /etc/quotas file

The /etc/quotas file consists of one or more entries, each entry specifying a default or explicit space or file quota limit for a qtree, group, or user.

The fields of a quota entry in the /etc/quotas file are

```
quota_target type[@/vol/dir/qtree_path] disk [files] [threshold]
[soft_disk] [soft_files]
```

The fields of an /etc/quotas file entry specify the following:

- *quota_target* specifies an explicit qtree, group, or user to which this quota is being applied. An asterisk (*) applies this quota as a default to all members of the type specified in this entry that do not have an explicit quota.
- *type* [*@/vol/dir/qtree_path*] specifies the type of entity (qtree, group, or user) to which this quota is being applied. If the type is user or group, this field can optionally restrict this user or group quota to a specific volume, directory, or qtree.
- *disk* is the disk space limit that this quota imposes on the qtree, group, user, or type in question.
- *files* (optional) is the limit on the number of files that this quota imposes on the qtree, group, or user in question.
- *threshold* (optional) is the disk space usage point at which warnings of approaching quota limits are issued.
- *soft_disk* (optional) is a soft quota space limit that, if exceeded, issues warnings rather than rejecting space requests.
- *soft_files* (optional) is a soft quota file limit that, if exceeded, issues warnings rather than rejecting file creation requests.

Note

For a detailed description of the above fields, see “Fields of the /etc/quotas file” on page 292.

Sample /etc/quotas file entries

The following sample quota entry assigns to user jsmith explicit limitations of 500 MB of disk space and 10,240 files in the rls volume and directory.

```
#Quota target  Type            Disk     Files    thold
#------------  -------------  ----     ------   -------
jsmith         user@/vol/rls   500m       10k   -
```

The following sample quota entry assigns to groups in the cad volume a default quota of 750 megabytes of disk space and 85,000 files per group. This quota applies to any group in the cad volume that does not have an explicit quota defined.

```
#Quota target  Type             Disk     Files    thold
#------------  -------------  ----     ------   -------
*              group@/vol/cad  750M       85K   -
```

Note

A line beginning with a pound sign (#) is considered a comment.

Each entry in the /etc/quotas file can extend to multiple lines, but the Files, Threshold, Soft Disk, and Soft Files fields must be on the same line as the Disk field. If they are not on the same line as the Disk field, they are ignored.

Order of entries

Entries in the /etc/quotas file can be in any order. After Data ONTAP receives a write request, it grants access only if the request meets the requirements specified by all /etc/quotas entries. If a quota target is affected by several /etc/quotas entries, the most restrictive entry applies.

Rules for a user or group quota

The following rules apply to a user or group quota:

- If you do not specify a path name to a volume or qtree to which the quota is applied, the quota takes effect in the root volume.
- You cannot impose restrictions on certain quota targets. For the following targets, you can specify quotas entries for tracking purposes only:
  - User with UID 0
  - Group with GID 0
  - BUILTIN\Administrators

- A file created by a member of the BUILTIN\Administrators group is owned by the BUILTIN\Administrators group, not by the member. When determining the amount of disk space or the number of files used by that user, Data ONTAP does not count the files that are owned by the BUILTIN\Administrators group.

Character coding of the /etc/quotas file

For information about character coding of the /etc/quotas file, see the *System Administration Guide*.

Fields of the /etc/quotas file

Quota Target field

The quota target specifies the user, group, or qtree to which you apply the quota. If the quota is a user or group quota, the same quota target can be in multiple /etc/quotas entries. If the quota is a tree quota, the quota target can be specified only once.

For a user quota: Data ONTAP applies a user quota to the user whose ID is specified in any format described in "How Data ONTAP identifies users for quotas" on page 284.

For a group quota: Data ONTAP applies a group quota to a GID, which you specify in the Quota Target field in any of these formats:

- The group name, such as publications
- The GID, such as 30
- A file or subdirectory whose GID matches the group, such as /vol/vol1/archive

Note

Specifying a file or directory name only enables Data ONTAP to obtain the GID. Data ONTAP does not apply quotas to that file or directory, or to the volume in which the file or directory resides.

For a tree quota: The quota target is the complete path name to an existing qtree (for example, /vol/vol0/home).

For default quotas: Use an asterisk (*) in the Quota Target field to specify a default quota. The quota is applied to the following users, groups, or qtrees:

- New users or groups that are created after the default entry takes effect. For example, if the maximum disk space for a default user quota is 500 MB, any new user can use up to 500 MB of disk space.
- Users or groups that are not explicitly mentioned in the /etc/quotas file. For example, if the maximum disk space for a default user quota is 500 MB, users for whom you have not specified a user quota in the /etc/quotas file can use up to 500 MB of disk space.

Type field

The Type field specifies the quota type, which can be

- User or group quotas, which specify the amount of disk space and the number of files that particular users and groups can own.
- Tree quotas, which specify the amount of disk space and the number of files that particular qtrees can contain.

For a user or group quota: The following table lists the possible values you can specify in the Type field, depending on the volume or the qtree to which the user or group quota is applied.

Quota type	Value in the Type field	Sample entry in the Type field
User quota in a volume	`user@/vol/`*`volume`*	`user@/vol/vol1`
User quota in a qtree	`user@/vol/`*`volume/qtree`*	`user@/vol/vol0/home`
Group quota in a volume	`group@/vol/`*`volume`*	`group@/vol/vol1`
Group quota in a qtree	`group@/vol/`*`volume/qtree`*	`group@/vol/vol0/home`

For a tree quota: The following table lists the values you can specify in the Type field, depending on whether the entry is an explicit tree quota or a default tree quota.

Entry	Value in the Type field
Explicit tree quota	`tree`
Default tree quota	`tree@/vol/`*`volume`* **Example:** `tree@/vol/vol0`

Disk field

The Disk field specifies the maximum amount of disk space that the quota target can use. The value in this field represents a hard limit that cannot be exceeded. The following list describes the rules for specifying a value in this field:

- K is equivalent to 1,024 bytes, M means 2^{20} bytes, and G means 2^{30} bytes.

Note

The Disk field is not case-sensitive. Therefore, you can use K, k, M, m, G, or g.

- The maximum value you can enter in the Disk field is 16 TB, or
  - 16,383G
  - 16,777,215M
  - 17,179,869,180K

Note

If you omit the K, M, or G, Data ONTAP assumes a default value of K.

- Your quota limit can be larger than the amount of disk space available in the volume. In this case, a warning message is printed to the console when quotas are initialized.
- The value cannot be specified in decimal notation.
- If you want to track the disk usage but do not want to impose a hard limit on disk usage, type a hyphen (-).
- Do not leave the Disk field blank. The value that follows the Type field is always assigned to the Disk field; thus, for example, Data ONTAP regards the following two quota file entries as equivalent:

```
#Quota Target         type           disk      files
 /export              tree           75K
 /export              tree                     75K
```

- If you do not specify disk space limits as a multiple of 4 KB, disk space fields can appear incorrect in quota reports. This happens because disk space fields are always rounded up to the nearest multiple of 4 KB to match disk space limits, which are translated into 4-KB disk blocks.

Files field

The Files field specifies the maximum number of files that the quota target can use. The value in this field represents a hard limit that cannot be exceeded. The following list describes the rules for specifying a value in this field:

- K is equivalent to 1,024, M means 2^{20}, and G means 2^{30}. You can omit the K, M, or G. For example, if you type 100, it means that the maximum number of files is 100.

Note

The Files field is not case-sensitive. Therefore, you can use K, k, M, m, G, or g.

- The maximum value you can enter in the Files field is 3GB, or
  - 4,294,967,295
  - 4,194,303K
  - 4,095M
  - 3G
- The value cannot be specified in decimal notation.
- If you want to track the number of files but do not want to impose a hard limit on the number of files that the quota target can use, type a hyphen (-). If the quota target is root, or if you specify 0 as the UID or GID, you must type a hyphen.
- A blank in this field means there is no restriction on the number of files that the quota target can use. If you leave this field blank, you cannot specify values for the Threshold, Soft Disk, or Soft Files fields.
- The Files field must be on the same line as the Disk field. Otherwise, the Files field is ignored.

Threshold field

The Threshold field specifies the disk space threshold. If a write causes the quota target to exceed the threshold, the write still succeeds, but a warning message is logged to the filer console and an SNMP trap is generated. Use the Threshold field to specify disk space threshold limits for CIFS.

The following list describes the rules for specifying a value in this field:

- The use of K, M, and G for the Threshold field is the same as for the Disk field.
- The maximum value you can enter in the Threshold field is 16 TB, or
  - 16,383G
  - 16,777,215M
  - 17,179,869,180K

Note

If you omit the K, M, or G, Data ONTAP assumes the default value of K.

- The value cannot be specified in decimal notation.
- The Threshold field must be on the same line as the Disk field. Otherwise, the Threshold field is ignored.

- If you do not want to specify a threshold limit on the amount of disk space the quota target can use, enter a hyphen (-) in this field or leave blank.

Note

Threshold fields can appear incorrect in quota reports if you do not specify threshold limits as multiples of 4 KB. This happens because threshold fields are always rounded up to the nearest multiple of 4 KB to match disk space limits, which are translated into 4-KB disk blocks.

Soft Disk field

The Soft Disk field specifies the amount of disk space that the quota target can use before a warning is issued. If the quota target exceeds the soft limit, a warning message is logged to the filer console and an SNMP trap is generated. When the soft disk limit is no longer being exceeded, another syslog message and SNMP trap are generated.

The following list describes the rules for specifying a value in this field:

- The use of K, M, and G for the Threshold field is the same as for the Disk field.
- The maximum value you can enter in the Soft Disk field is 16 TB, or
  - 16,383G
  - 16,777,215M
  - 17,179,869,180K
- The value cannot be specified in decimal notation.
- If you do not want to specify a soft limit on the amount of disk space that the quota target can use, type a hyphen (-) in this field (or leave this field blank if no value for the Soft Files field follows).
- The Soft Disk field must be on the same line as the Disk field. Otherwise, the Soft Disk field is ignored.

Note

Disk space fields can appear incorrect in quota reports if you do not specify disk space limits as multiples of 4 KB. This happens because disk space fields are always rounded up to the nearest multiple of 4 KB to match disk space limits, which are translated into 4-KB disk blocks.

Soft Files field

The Soft Files field specifies the number of files that the quota target can use before a warning is issued. If the quota target exceeds the soft limit, a warning message is logged to the filer console and an SNMP trap is generated. When the soft files limit is no longer being exceeded, another syslog message and SNMP trap are generated.

The following list describes the rules for specifying a value in this field.

- The format of the Soft Files field is the same as the format of the Files field.
- The maximum value you can enter in the Soft Files field is 4,294,967,295.
- The value cannot be specified in decimal notation.
- If you do not want to specify a soft limit on the number of files that the quota target can use, type a hyphen (-) in this field or leave the field blank.
- The Soft Files field must be on the same line as the Disk field. Otherwise, the Soft Files field is ignored.

Sample quota entries

Explicit quota examples

The following list contains examples of explicit quotas:

- `jsmith  user@/vol/rls  500M  10K`

  The user named jsmith can use 500 MB of disk space and 10,240 files in the rls volume.

- `jsmith,corp\jsmith,engineering\"john smith",`
  `S-1-5-32-544  user@/vol/rls  500M  10K`

  This user, represented by four IDs, can use 500 MB of disk space and 10,240 files in the rls volume.

- `writers  group@/vol/cad/proj1  150M`

  The writers group can use 150 MB of disk space and an unlimited number of files in the /vol/cad/proj1 qtree.

- `/vol/cad/proj1  tree  750M  75K`

  The proj1 qtree in the cad volume can use 750 MB of disk space and 76,800 files.

Tracking quota examples

The following list contains examples of tracking quotas:

- `root  user@/vol/rls  -  -`

  Data ONTAP tracks but does not limit the amount of disk space and the number of files in the rls volume owned by root.

- `builtin\administrators  user@/vol/rls  -  -`

  Data ONTAP tracks but does not limit the amount of disk space and the number of files in the rls volume owned by or created by members of BUILTIN\Administrators.

- `/vol/cad/proj1  tree  -  -`

  Data ONTAP tracks but does not limit the amount of disk space and the number of files for the proj1 qtree in the cad volume.

Default quota examples

The following list contains examples of default quotas:

- `*  user@/vol/cad  50M  15K`

  Any user not explicitly listed in the quota file can use 50 MB of disk space and 15,360 files in the cad volume.
- `*  group@/vol/cad  750M  85K`

  Any group not explicitly listed in the quota file can use 750 MB of disk space and 87,040 files in the cad volume.
- `*  tree@vol/cad  75M`

  Any qtree in the cad volume that is not explicitly listed in the quota file can use 75 MB of disk space and an unlimited number of files.

Default tracking quota example

Default tracking quotas enable you to create default quotas that do not enforce any resource limits. This is helpful when you want to use the `quota resize` command when you modify your /etc/quotas file, but you do not want to apply resource limits with your default quotas. Default tracking quotas are created per-volume, as shown in the following example:

```
#Quota Target         type                  disk      files
*                     user@/vol/vol1          -         -
*                     group@/vol/vol1         -         -
*                     tree@/vol/vol1          -         -
```

Sample quota file and explanation

The following sample /etc/quotas file contains default quotas and explicit quotas:

```
#Quota Target         type                  disk      files
*                     user@/vol/cad          50M      15K
*                     group@/vol/cad        750M      85K
*                     tree@/vol/cad         100M      75K
jdoe                  user@/vol/cad/proj1   100M      75K
msmith                user@/vol/cad          75M      75K
msmith                user@/vol/cad/proj1    75M      75K
```

The following list explains the effects of these /etc/quotas entries:

- Any user not otherwise mentioned in this file can use 50 MB of disk space and 15,360 files in the cad volume.
- Any group not otherwise mentioned in this file can use 750 MB of disk space and 87,040 files in the cad volume.
- Any qtree in the cad volume not otherwise mentioned in this file can use 100 MB of disk space and 76,800 files.

- If a qtree is created in the cad volume (for example, a qtree named /vol/cad/proj2), Data ONTAP enforces a derived default user quota and a derived default group quota that have the same effect as these quota entries:

  ```
  *                 user@/vol/cad/proj2     50M      15K
  *                 group@/vol/cad/proj2    750M     85K
  ```

- If a qtree is created in the cad volume (for example, a qtree named /vol/cad/proj2), Data ONTAP tracks the disk space and number of files owned by UID 0 and GID 0 in the /vol/cad/proj2 qtree. This is due to this quota file entry:

  ```
  *                 tree@/vol/cad           100M     75K
  ```

- A user named msmith can use 75 MB of disk space and 76,800 files in the cad volume because an explicit quota for this user exists in the /etc/quotas file, overriding the default limit of 50 MB of disk space and 15,360 files.
- By giving jdoe and msmith 100 MB and 75 MB explicit quotas for the proj1 qtree, which has a tree quota of 100MB, that qtree becomes oversubscribed. This means that the qtree could run out of space before the user quotas are exhausted.

  Quota oversubscription is supported; however, a warning is printed alerting you to the oversubscription.

How conflicting quotas are resolved

When more than one quota is in effect, the most restrictive quota is applied. Consider the following example /etc/quota file:

```
*                   tree@/vol/cad          100M      75K
jdoe                user@/vol/cad/proj1    750M      75K
```

Because the jdoe user has a disk quota of 750 MB in the proj1 qtree, you might expect that to be the limit applied in that qtree. But the proj1 qtree has a tree quota of 100 MB, because of the first line in the quota file. So jdoe will not be able to write more than 100 MB to the qtree. If other users have already written to the proj1 qtree, the limit would be reached even sooner.

To remedy this situation, you can create an explicit tree quota for the proj1 qtree, as shown in this example:

```
*                   tree@/vol/cad          100M      75K
/vol/cad/proj1      tree                   800M      75K
jdoe                user@/vol/cad/proj1    750M      75K
```

Now the jdoe user is no longer restricted by the default tree quota and can use the entire 750 MB of the user quota in the proj1 qtree.

Special entries for mapping users

Special entries in the /etc/quotas file

The /etc/quotas file supports two special entries whose formats are different from the entries described in "Fields of the /etc/quotas file" on page 292. These special entries enable you to quickly add Windows IDs to the /etc/quotas file. If you use these entries, you can avoid typing individual Windows IDs.

These special entries are

- QUOTA_TARGET_DOMAIN
- QUOTA_PERFORM_USER_MAPPING

Note

If you add or remove these entries from the /etc/quotas file, you must perform a full quota reinitialization for your changes to take effect. A `quota resize` command is not sufficient. For more information about quota reinitialization, see "Modifying quotas" on page 309.

Special entry for changing UNIX names to Windows names

The QUOTA_TARGET_DOMAIN entry enables you to change UNIX names to Windows names in the Quota Target field. Use this entry if both of the following conditions apply:

- The /etc/quotas file contains user quotas with UNIX names.
- The quota targets you want to change have identical UNIX and Windows names. For example, a user whose UNIX name is jsmith also has a Windows name of jsmith.

Format: The following is the format of the QUOTA_TARGET_DOMAIN entry:

```
QUOTA_TARGET_DOMAIN domain_name
```

Effect: For each user quota, Data ONTAP adds the specified domain name as a prefix to the user name. Data ONTAP stops adding the prefix when it reaches the end of the /etc/quotas file or another QUOTA_TARGET_DOMAIN entry without a domain name.

Example: The following example illustrates the use of the QUOTA_TARGET_DOMAIN entry:

```
QUOTA_TARGET_DOMAIN corp
roberts     user@/vol/rls          900M      30K
smith       user@/vol/rls          900M      30K
QUOTA_TARGET_DOMAIN engineering
daly        user@/vol/rls          900M      30K
thomas      user@/vol/rls          900M      30K
QUOTA_TARGET_DOMAIN
stevens     user@/vol/rls          900M      30K
```

Explanation of example: The string corp\ is added as a prefix to the user names of the first two entries. The string engineering\ is added as a prefix to the user names of the third and fourth entries. The last entry is unaffected by the QUOTA_TARGET_DOMAIN entry. The following entries produce the same effects:

```
corp\roberts        user@/vol/rls          900M      30K
corp\smith          user@/vol/rls          900M      30K
engineering\daly    user@/vol/rls          900M      30K
engineering\thomas  user@/vol/rls          900M      30K
stevens             user@/vol/rls          900M      30K
```

Special entry for mapping names

The QUOTA_PERFORM_USER_MAPPING entry enables you to map UNIX names to Windows names or vice versa. Use this entry if both of the following conditions apply:

- There is a one-to-one correspondence between UNIX names and Windows names.
- You want to apply the same quota to the user whether the user uses the UNIX name or the Windows name.

Note

The QUOTA_PERFORM_USER_MAPPING entry does not work if the QUOTA_TARGET_DOMAIN entry is present.

How names are mapped: Data ONTAP consults the /etc/usermap.cfg file to map the user names. For more information about how Data ONTAP uses the usermap.cfg file, see the *File Access and Protocols Management Guide*.

Format: The QUOTA_PERFORM_USER_MAPPING entry has the following format:

```
QUOTA_PERFORM_USER_MAPPING [ on | off ]
```

Data ONTAP maps the user names in the Quota Target fields of all entries following the `QUOTA_PERFORM_USER_MAPPING on` entry. It stops mapping when it reaches the end of the /etc/quotas file or when it reaches a `QUOTA_PERFORM_USER_MAPPING off` entry.

Note

If a default user quota entry is encountered after the QUOTA_PERFORM_USER_MAPPING directive, any user quotas derived from that default quota are also mapped.

Example: The following example illustrates the use of the QUOTA_PERFORM_USER_MAPPING entry:

```
QUOTA_PERFORM_USER_MAPPING on
roberts         user@/vol/rls          900M      30K
corp\stevens    user@/vol/rls          900M      30K
QUOTA_PERFORM_USER_MAPPING off
```

Explanation of example: If the /etc/usermap.cfg file maps roberts to corp\jroberts, the first quota entry applies to the user whose UNIX name is roberts and whose Windows name is corp\jroberts. A file owned by a user with either user name is subject to the restriction of this quota entry.

If the usermap.cfg file maps corp\stevens to cws, the second quota entry applies to the user whose Windows name is corp\stevens and whose UNIX name is cws. A file owned by a user with either user name is subject to the restriction of this quota entry.

The following entries produce the same effects:

```
roberts,corp\jroberts        user@/vol/rls          900M      30K
corp\stevens,cws             user@/vol/rls          900M      30K
```

Importance of one-to-one mapping: If the name mapping is not one-to-one, the QUOTA_PERFORM_USER_MAPPING entry produces confusing results, as illustrated in the following examples.

Example of multiple Windows names for one UNIX name: Suppose the /etc/usermap.cfg file contains the following entries:

```
domain1\user1 => unixuser1
domain2\user2 => unixuser1
```

Data ONTAP displays a warning message if the /etc/quotas file contains the following entries:

```
QUOTA_PERFORM_USER_MAPPING on
domain1\user1 user 1M
domain2\user2 user 1M
```

The /etc/quotas file effectively contains two entries for unixuser1. Therefore, the second entry is treated as a duplicate entry and is ignored.

Example of wildcard entries in usermap.cfg: Confusion can result if the following conditions exist:

- The /etc/usermap.cfg file contains the following entry:

  ```
  *\* *
  ```

- The /etc/quotas file contains the following entries:

  ```
  QUOTA_PERFORM_USER_MAPPING on
  unixuser2 user 1M
  ```

Problems arise because Data ONTAP tries to locate unixuser2 in one of the trusted domains. Because Data ONTAP searches domains in an unspecified order, unless the order is specified by the `cifs.search_domains` option, the result becomes unpredictable.

What to do after you change usermap.cfg: If you make changes to the /etc/usermap.cfg file, you must turn quotas off and then turn quotas back on for the changes to take effect. For more information about turning quotas on and off, see “Activating or reinitializing quotas” on page 306.

How disk space owned by default users is counted

Disk space used by the default UNIX user

For a Windows name that does not map to a specific UNIX name, Data ONTAP uses the default UNIX name defined by the `wafl.default_unix_user` option when calculating disk space. Files owned by the Windows user without a specific UNIX name are counted against the default UNIX user name if either of the following conditions applies:

- The files are in qtrees with UNIX security style.
- The files do not have ACLs in qtrees with mixed security style.

Disk space used by the default Windows user

For a UNIX name that does not map to a specific Windows name, Data ONTAP uses the default Windows name defined by the `wafl.default_nt_user` option when calculating disk space. Files owned by the UNIX user without a specific Windows name are counted against the default Windows user name if the files have ACLs in qtrees with NTFS security style or mixed security style.

Activating or reinitializing quotas

About activating or reinitializing quotas

You use the `quota on` command to activate or reinitialize quotas. The following list outlines some facts you should know about activating or reinitializing quotas:

- You activate or reinitialize quotas for only one volume at a time.
- In Data ONTAP 7.0 and later, it is no longer a requirement for activating quotas that your /etc/quotas file be free of any errors. Invalid entries are reported and skipped. If the /etc/quotas file contains any valid entries, quotas are activated.
- Reinitialization causes the quota file to be scanned and all quotas for that volume to be recalculated.
- Changes to the /etc/quotas file do not take effect until either quotas are reinitialized or the `quota resize` command is issued.
- Quota reinitialization can take some time, during which NetApp system data is available, but quotas are not enforced for the specified volume.
- Quota reinitialization is performed asynchronously by default; other commands can be performed while the reinitialization is proceeding in the background.

Note

This means that errors or warnings from the reinitialization process could be interspersed with the output from other commands.

- Quota reinitialization can be invoked synchronously with the `-w` option; this is useful if you are reinitializing from a script.
- Errors and warnings from the reinitialization process are logged to the console as well as to /etc/messages.

Note

For more information about when to use the `quota resize` command versus the `quota on` command after changing the quota file, see “Modifying quotas” on page 309.

CIFS requirement for activating quotas

If the /etc/quotas file contains user quotas that use Windows IDs as targets, CIFS must be running before you can activate or reinitialize quotas.

Quota initialization terminated by upgrade

In previous versions of Data ONTAP, if an upgrade was initiated while a quota initialization was in progress, the initialization completed after the filer came back online. In Data ONTAP 7.0, any quota initialization running when the filer is upgraded is terminated and must be manually restarted from the beginning. For this reason, NetApp recommends that you allow any running quota initialization to complete before upgrading your filer.

Activating quotas

To activate quotas, complete the following step.

Step	Action
1	Enter the following command: **`quota on [-w] vol_name`** The `-w` option causes the command to return only after the entire /etc/quotas file has been scanned (synchronous mode). This is useful when activating quotas from a script. **Example:** The following example turns on quotas on a volume named cad: `quota on cad`

Reinitializing quotas

To reinitialize quotas, complete the following steps.

Step	Action
1	If quotas are already on for the volume you want to reinitialize quotas on, enter the following command: **`quota off vol_name`**
2	Enter the following command: **`quota on vol_name`**

Deactivating quotas

To deactivate quotas, complete the following step.

Step	Action
1	Enter the following command: **quota off *vol_name*** **Example:** The following example turns off quotas on a volume named cad: `quota off cad` **Note** If a quota initialization is almost complete, the `quota off` command can fail. If this happens, retry the command after a minute or two.

Canceling quota initialization

To cancel a quota initialization that is in progress, complete the following step.

Step	Action
1	Enter the following command: **quota off *vol_name*** **Note** If a quota initialization is almost complete, the `quota off` command can fail. In this case, the initialization scan is already complete.

Modifying quotas

About modifying quotas

When you want to change how quotas are being tracked on your filer, you first need to make the required change to your /etc/quota file. Then, you need to request Data ONTAP to read the /etc/quota file again and incorporate the changes. You can do this using one of the two following methods:

- Resize quotas

  Resizing quotas is faster than a full reinitialization; however, some quota file changes may not be reflected.

- Reinitialize quotas

  Performing a full quota reinitialization reads and recalculates the entire quota file. This may take some time, but all quota file changes are guaranteed to be reflected after the initialization is complete.

Note

Your filer functions normally while quotas are being initialized; however, quotas remain off until the initialization is complete.

When you can use resizing

Because quota resizing is faster than quota initialization, you should use resizing whenever possible. You can use quota resizing for the following types of changes to the /etc/quota file:

- You changed an existing quota file entry, including adding or removing fields.
- You added a quota file entry for a quota target that was already covered by a default or default tracking quota.
- You deleted an entry from your /etc/quota file for which a default or default tracking quota entry is specified.

Note

After you have made extensive changes to the /etc/quota file, NetApp recommends that you perform a full reinitialization to ensure that all of the changes become effective.

Resizing example 1: Consider the following sample /etc/quota file:

```
#Quota Target   Type                Disk    Files
*               user@/vol/cad        50M    15K
*               group@/vol/cad      750M    85K
*               tree@vol/cad          -      -
jdoe            user@/vol/cad/      100M    75K
kbuck           user@/vol/cad/      100M    75K
```

Suppose you make the following changes:

- Increase the number of files for the default user target.
- Added a new user quota for a new user that needs more than the default user quota.
- Deleted the kbuck user's explicit quota entry; the kbuck user now only needs the default quota limits.

These changes result in the following /etc/quota file:

```
#Quota Target   Type                Disk    Files
*               user@/vol/cad        50M    25K
*               group@/vol/cad      750M    85K
*               tree@vol/cad          -      -
jdoe            user@/vol/cad/      100M    75K
bambi           user@/vol/cad/      100M    75K
```

All of these changes can be made effective using the `quota resize` command; a full quota reinitialization is not necessary.

Resizing example 2: Your quotas file did not contain the default tracking tree quota, and you want to add a tree quota to the sample quota file, resulting in this /etc/quota file:

```
#Quota Target   Type                Disk    Files
*               user@/vol/cad        50M    25K
*               group@/vol/cad      750M    85K
jdoe            user@/vol/cad/      100M    75K
bambi           user@/vol/cad/      100M    75K
/vol/cad/proj1  tree                500M   100K
```

In this case, using the `quota resize` command does *not* cause the newly added entry to be effective, because there is no default entry for tree quotas already in effect. A full quota initialization is required.

Note

If you use the `resize` command and the /etc/quota file contains changes that will not be reflected, Data ONTAP issues a warning.

You can determine from the quota report whether your filer is tracking disk usage for a particular user, group, or qtree. A quota in the quota report indicates that the filer is tracking the disk space and the number of files owned by the quota target. For more information about quota reports, see "Understanding quota reports" on page 318.

Resizing quotas

To resize quotas, complete the following step.

Step	Action
1	Enter the following command: **`quota resize`** ***`vol_name`*** *vol_name* is the name of the volume you want to resize quotas for.

Deleting quotas

About quota deletion

You can remove quota restrictions for a quota target in two ways:

- Delete the /etc/quotas entry pertaining to the quota target.

  If you have a default or default tracking quota entry for the target type you deleted, you can use the `quota resize` command to update your quotas. Otherwise, you must reinitialize quotas.
- Change the /etc/quotas entry so that there is no restriction on the amount of disk space or the number of files owned by the quota target. After the change, Data ONTAP continues to keep track of the disk space and the number of files owned by the quota target but stops imposing the restrictions on the quota target. The procedure for removing quota restrictions in this way is the same as that for resizing an existing quota.

  You can use the `quota resize` command after making this kind of modification to the quotas file.

Deleting a quota by removing restrictions

To delete a quota by removing the resource restrictions for the specified target, complete the following steps.

Step	Action

1	Open the /etc/quotas file and edit the quotas file entry for the specified target so that the quota entry becomes a tracking quota.
	Example: Your quota file contains the following entry for the jdoe user:
	`jdoe  user@/vol/cad/  100M  75K`
	To remove the restrictions on jdoe, edit the entry as follows:
	`jdoe  user@/vol/cad/  -  -`

2	Enter the following command to update quotas: **`quota resize`** ***`vol_name`***

Deleting a quota by removing the quota file entry

To delete a quota by removing the quota file entry for the specified target, complete the following steps.

Step	Action	
1	Open the /etc/quotas file and remove the entry for the quota you want to delete.	
2	**If…**	**Then…**
	You have a default or default tracking quotas in place for users, groups and qtrees	Enter the following command to update quotas: **`quota resize`** ***`vol_name`***
	Otherwise	Enter the following commands to reinitialize quotas: **`quota off`** ***`vol_name`*** **`quota on`** ***`vol_name`***

Turning quota message logging on or off

About turning quota message logging on or off

You can turn quota message logging on or off for a single volume or for all volumes. You can optionally specify a time interval during which quota messages will not be logged.

Turning quota message logging on

To turn quota message logging on, complete the following step.

Step	Action
1	Enter the following command: **`quota logmsg on [ interval ] [ -v vol_name | all ]`** *interval* is the time period during which quota message logging is disabled. The interval is a number followed by `d`, `h`, or `m` for days, hours, and minutes, respectively. Quota messages are logged after the end of each interval. If no interval is specified, Data ONTAP logs quota messages every 60 minutes. For continuous logging, specify `0m` for the interval. `-v` *vol_name* specifies a volume name. `all` applies the interval to all volumes in the filer. **Note** If you specify a short interval, less than five minutes, quota messages might not be logged exactly at the specified rate because Data ONTAP buffers quota messages before logging them.

Turning quota message logging off

To turn quota message logging off, complete the following step.

Step	Action
1	Enter the following command: **`quota logmsg off`**

Displaying settings for quota message logging

To display the current settings for quota message logging, complete the following step.

Step	Action
1	Enter the following command: `quota logmsg`

Effects of qtree changes on quotas

Effect of deleting a qtree on tree quotas

When you delete a qtree, all quotas that are applicable to that qtree, whether they are explicit or derived, are automatically deleted.

If you create a new qtree with the same name as the one you deleted, the quotas previously applied to the deleted qtree are not applied automatically to the new qtree. If a default tree quota exists, Data ONTAP creates new derived quotas for the new qtree. However, explicit quotas in the /etc/quotas file do not apply until you reinitialize quotas.

Effect of renaming a qtree on tree quotas

When you rename a qtree, Data ONTAP keeps the same ID for the tree. As a result, all quotas applicable to the qtree, whether they are explicit or derived, continue to be applicable.

Effects of changing qtree security style on user quota usages

Because ACLs apply in qtrees using NTFS or mixed security style but not in qtrees using UNIX security style, changing the security style of a qtree through the `qtree security` command might affect how a UNIX or Windows user's quota usages for that qtree are calculated.

Example: If NTFS security is in effect on qtree A and an ACL gives Windows user Windows/joe ownership of a 5-MB file, then user Windows/joe is charged 5 MB of quota usage on qtree A.

If the security style of qtree A is changed to UNIX, and Windows user Windows/joe is default mapped to UNIX user joe, the ACL that charged 5 MB of diskspace against the quota of Windows/joe is ignored when calculating the quota usage of UNIX user joe.

Caution

To make sure quota usages for both UNIX and Windows users are properly calculated after you use the `qtree security` command to change the security style, turn quotas for the volume containing that qtree off and then back on again using the `quota off` *vol_name* and `quota on` *vol_name* commands.

If you change the security style from UNIX to either mixed or NTFS, previously hidden ACLs become visible, any ACLs that were ignored become effective again, and the NFS user information is ignored. If no ACL existed before, then the NFS information is used in the quota calculation.

Note

Only UNIX group quotas apply to qtrees. Changing the security style of a qtree, therefore, does not affect the quota usages that groups are subject to.

Understanding quota reports

About this section

This section provides information about quota reports.

Detailed information

The following sections provide detailed information about quota reports:

- "Types of quota reports" on page 319
- "Overview of the quota report format" on page 320
- "Quota report formats" on page 322
- "Displaying a quota report" on page 326

Types of quota reports

Types of quota reports

You can display these types of quota reports:

- A quota report for all volumes that have quotas turned on. It contains the following types of information:
  - Default quota information, which is the same information as that in the /etc/quotas file
  - Current disk space and the number of files owned by a user, group, or qtree that has an explicit quota in the /etc/quotas file
  - Current disk space and the number of files owned by a user, group, or qtree that is the quota target of a derived quota, if the user, group, or qtree currently uses some disk space
- A quota report for a specified path name. It contains information about all the quotas that apply to the specified path name.

  For example, in the quota report for the /vol/cad/specs path name, you can see the quotas to which the disk space used by the /vol/cad/specs path name is charged. If a user quota exists for the owner of the /vol/cad/specs path name and a group quota exists for the cad volume, both quotas appear in the quota report.

Overview of the quota report format

Contents of the quota report

The following table lists the fields displayed in the quota report and the information they contain.

Heading	Information
Type	Quota type: user, group, or tree.
ID	User ID, UNIX group name, qtree name. If the quota is a default quota, the value in this field is an asterisk.
Volume	Volume to which the quota is applied.
Tree	Qtree to which the quota is applied.
K-Bytes Used	Current amount of disk space used by the quota target. If the quota is a default quota, the value in this field is 0.
Limit	Maximum amount of disk space that can be used by the quota target (Disk field).
S-Limit	Maximum amount of disk space that can be used by the quota target before a warning is issued (Soft Disk field). This column is displayed only when you use the `-s` option for the `quota report` command.
T-hold	Disk space threshold (Threshold field). This column is displayed only when you use the `-t` option for the `quota report` command.
Files Used	Current number of files used by the quota target. If the quota is a default quota, the value in this field is 0. If a soft files limit is specified for the quota target, you can also display the soft files limit in this field.

Heading	Information
Limit	Maximum number of files allowed for the quota target (Files field).
S-Limit	Maximum number of files that can be used by the quota target before a warning is issued (Soft Files field). This column is displayed only when you use the `-s` option for the `quota report` command.
Vfiler	Displays the name of the vFiler for this quota entry. This column is displayed only when you use the `-v` option for the `quota report` command, which is available only on filers that have MultiStore licensed.
Quota Specifier	For an explicit quota, it shows how the quota target is specified in the /etc/quotas file. For a derived quota, the field is blank.

Quota report formats

Available report formats

Quota reports are available in these formats:

- A default format generated by the `quota report` command

  For more information, see "Default format" on page 323.
- Target IDs displayed in numeric form using the `quota report -q` command

  For more information, see "Report format with quota report -q" on page 324.
- Soft limits listed using the `quota report -s` command
- Threshold values listed using the `quota report -t` command
- vFiler names included using the `quota report -v` command

  This option is valid only if MultiStore is licensed.
- Two enhanced formats for quota targets with multiple IDs:
  - IDs listed on different lines using the `quota report -u` command

    For more information, see "Report format with quota report -u" on page 324.
  - IDs listed in a comma separated list using the `quota report -x` command

    For more information, see "Report format with quota report -x" on page 325.

Factors affecting the contents of the fields

The information contained in the ID and Quota Specifier fields can vary according to these factors:

- Type of user—UNIX or Windows—to which a quota applies
- The specific command used to generate the quota report

Contents of the ID field

In general, the ID field of the quota report displays a user name instead of a UID or SID; however, the following exceptions apply:

- For a quota with a UNIX user as the target, the ID field shows the UID instead of a name if no user name for the UID is found in the password database, or if you specifically request the UID by including the `-q` option in the `quota reports` command.

- For a quota with a Windows user as the target, the ID field shows the SID instead of a name if either of the following conditions applies:
  - The SID is specified as a quota target and the SID no longer corresponds to a user name.
  - The filer cannot find an entry for the SID in the SID-to-name map cache and cannot connect to the domain controller to ascertain the user name for the SID when it generates the quota report.

Default format

The `quota report` command without options generates the default format for the ID and Quota Specifier fields.

The ID field: If a quota target contains only one ID, the ID field displays that ID. Otherwise, the ID field displays one of the IDs from the list.

The ID field displays information in the following formats:

- For a Windows name, the first seven characters of the user name with a preceding backslash are displayed. The domain name is omitted.
- For a SID, the last eight characters are displayed.

The Quota Specifier field: The Quota Specifier field displays an ID that matches the one in the ID field. The ID is displayed the same way the quota target is specified in the /etc/quotas file.

Examples: The following table shows what is displayed in the ID and Quota Specifier fields based on the quota target in the /etc/quotas file.

Quota target in the /etc/quotas file	ID field of the quota report	Quota Specifier field of the quota report
CORP\john_smith	\john_sm	CORP\john_smith
CORP\john_smith,NT\js	\john_sm or \js	CORP\john_smith or NT\js
S-1-5-32-544	5-32-544	S-1-5-32-544

Report format with quota report -q

The `quota report -q` command displays the quota target's UNIX UID or GID in numeric form. No lookup of the name associated with the target ID is done.

For Windows IDs, the textual form of the SID is displayed.

UNIX UIDs and GIDs are displayed as numbers. Windows SIDs are displayed as text.

Report format with quota report -s

The format of the report generated using the `quota report -s` command is the same as the default format, except that the soft limit columns are included.

Report format with quota report -t

The format of the report generated using the `quota report -t` command is the same as the default format, except that the threshold column is included.

Report format with quota report -v

The format of the report generated using the `quota report -v` command is the same as the default format, except that the Vfiler column is included. This format is available only if MultiStore is licensed.

Report format with quota report -u

The `quota report -u` command is useful if you have quota targets that have multiple IDs. It provides more information in the ID and Quota Specifier fields than the default format.

If a quota target consists of multiple IDs, the first ID is listed on the first line of the quota report for that entry. The other IDs are listed on the lines following the first line, one ID per line. Each ID is followed by its original quota specifier, if any. Without this option, only one ID is displayed for quota targets with multiple IDs.

Note

You cannot combine the `-u` and `-x` options.

The ID field: The ID field displays all the IDs listed in the quota target of a user quota in the following format:

- On the first line, the format is the same as the default format.
- Each additional name in the quota target is displayed on a separate line in its entirety.

The Quota Specifier field: The Quota Specifier field displays the same list of IDs as specified in the quota target.

Example: The following table shows what is displayed in the ID and Quota Specifier fields based on the quota target in the /etc/quotas file. In this example, the SID maps to the user name NT\js.

Quota target in /etc/quotas	ID field of the quota report	Quota Specifier field of the quota report
CORP\john_smith,S-1-5-21-123456-7890-1234-1166	\john_sm NT\js	CORP\john_smith,S-1-5-21-123456-7890-1234-1166

Report format with quota report -x

The `quota report -x` command report format is similar to the report displayed by the `quota report -u` command, except that `quota report -x` displays all the quota target's IDs on the first line of that quota target's entry, as a comma separated list. The threshold column is included.

Note

You cannot combine the `-x` and `-u` options.

Displaying a quota report

Displaying a quota report for all quotas

To display a quota report for all quotas, complete the following step.

Step	Action
1	Enter the following command: **`quota report [-q] [-s] [-t] [-v] [-u|-x]`** For complete information on the `quota report` options, see "Quota report formats" on page 322.

Displaying a quota report for a specified path name

To display a quota report for a specified path name, complete the following step.

Step	Action
1	Enter the following command: **`quota report [-s] [ -u | -x ] [ -t ] [-q] path_name`** *path_name* is a complete path name to a file, directory, or volume, such as /vol/vol0/etc. For complete information on the `quota report` options, see "Quota report formats" on page 322.

SnapLock Management 9

About this chapter

This chapter describes how to use SnapLock volumes and aggregates to provide WORM (write-once-read-many) storage.

Topics in this chapter

This chapter discusses the following topics:

- "About SnapLock" on page 328
- "Creating SnapLock volumes" on page 330
- "Managing the compliance clock" on page 332
- "Setting volume retention periods" on page 334
- "Destroying SnapLock volumes and aggregates" on page 337
- "Managing WORM data" on page 339

About SnapLock

What SnapLock is

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

SnapLock is available in two forms: SnapLock Compliance and SnapLock Enterprise.

SnapLock Compliance: Provides WORM protection of files while also restricting the storage administrator's ability to perform any operations that might modify or erase retained WORM records. SnapLock Compliance should be used in strictly regulated environments that require information to be retained for specified lengths of time, such as those governed by SEC Rule 17a-4.

SnapLock Enterprise: Provides WORM protection of files, but uses a trusted administrator model of operation that allows the storage administrator to manage the system with very few restrictions. For example, SnapLock Enterprise allows the administrator to perform operations, such as destroying SnapLock volumes, that might result in the loss of data.

Note

SnapLock Enterprise should not be used in strictly regulated environments.

How SnapLock works

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

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

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

In Data ONTAP 7.0 and later versions, WORM files can be deleted after their retention date. The retention date on a WORM file is set when the file is committed to WORM state, but can be extended at any time. The retention period can never be shortened for any WORM file.

Licensing SnapLock functionality

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

- SnapLock Compliance

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

- SnapLock Enterprise

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

AutoSupport with SnapLock

If AutoSupport is enabled, the filer sends AutoSupport messages to NetApp Technical Support. These messages include event and log-level descriptions. SnapLock volume state and options are included in AutoSupport output.

Replicating SnapLock volumes

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

Creating SnapLock volumes

SnapLock is an attribute of the containing aggregate

Although this guide uses the term “SnapLock volume” to describe volumes that contain WORM data, in fact SnapLock is an attribute of the volume’s containing aggregate. Because traditional volumes have a one-to-one relationship with their containing aggregate, you create traditional SnapLock volumes much as you would a standard traditional volume. To create SnapLock FlexVol volumes, you must first create a SnapLock aggregate. Every FlexVol created in that SnapLock aggregate is, by definition, a SnapLock volume.

Creating SnapLock traditional volumes

SnapLock traditional volumes are created in the same way a standard traditional volume is created, except that you use the `-L` parameter with the `vol create` command.

For more information about the `vol create` command, see “Creating traditional volumes” on page 182.

Verifying volume status

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

```
sys1> vol status

      Volume State      Status            Options
        vol0 online     raid4, trad       root
    wormvol online      raid4, trad       no_atime_update=on,
                                          snaplock_compliance
```

Creating SnapLock aggregates

SnapLock aggregates are created in the same way a standard aggregate is created, except that you use the `-L` parameter with the `aggr create` command.

For more information about the `aggr create` command, see “Creating aggregates” on page 157.

Verifying aggregate status

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

```
sys1> aggr status

       Aggr State      Status            Options
       vol0 online     raid4, trad       root
  wormaggr online      raid4, aggr       snaplock_compliance
```

SnapLock write_verify option

Data ONTAP provides a write verification option for SnapLock Compliance volumes: `snaplock.compliance.write_verify`. When this option is enabled, an immediate read verification occurs after every disk write, providing an additional level of data integrity.

Note

The SnapLock write verification option provides negligible benefit beyond the advanced, high-performance data protection and integrity features already provided by NVRAM, checksums, RAID scrubs, media scans, and double-parity RAID. SnapLock write verification should be used where the interpretation of regulations requires that each write to the disk media be immediately read back and verified for integrity.

SnapLock write verification comes at a performance cost and may affect data throughput on SnapLock Compliance volumes.

Managing the compliance clock

SnapLock Compliance requirements to enforce WORM retention

SnapLock Compliance meets the following requirements needed to enforce WORM data retention:

- Secure time base—ensures that retained data cannot be deleted prematurely by changing the regular clock of the storage system
- Synchronized time source—provides a time source for general operation that is synchronized to a common reference time used inside your data center

How SnapLock Compliance meets the requirements

SnapLock Compliance meets the requirements by using a secure compliance clock. The compliance clock is implemented in software and runs independently of the system clock. Although running independently, the compliance clock tracks the regular system clock and remains very accurate with respect to the system clock.

Initializing the compliance clock

To initialize the compliance clock, complete the following steps.

Caution

The compliance clock can be initialized *only once* for the system. You should exercise extreme care when setting the compliance clock to ensure that you set the compliance clock time correctly.

Step	Action
1	Ensure that the system time and time zone are set correctly.
2	Initialize the compliance clock using the following command: `date -c initialize` **Result:** The system prompts you to confirm the current local time and that you want to initialize the compliance clock.
3	Confirm that the system clock is correct and that you want to initialize the compliance clock.

Example: `filer> date -c initialize`

```
*** WARNING: YOU ARE INITIALIZING THE SECURE COMPLIANCE CLOCK ***

You are about to initialize the secure Compliance Clock of this
system to the current value of the system clock. This procedure
can be performed ONLY ONCE on this system so you should ensure
that the system time is set correctly before proceeding.

The current local system time is: Wed Feb 4 23:38:58 GMT 2004

Is the current local system time correct? y
Are you REALLY sure you want initialize the Compliance Clock? y

Compliance Clock: Wed Feb 4 23:39:27 GMT 2004
```

Viewing the compliance clock time

To view the compliance clock time, complete the following step.

Step	Action

1	Enter the command:
	`date -c`
	Example:
	`date -c`
	`Compliance Clock: Wed Feb 4 23:42:39 GMT 2004`

Setting volume retention periods

When you should set the retention periods

You should set the retention periods after creating the SnapLock volume and before using the SnapLock volume. Setting the options at this time ensures that the SnapLock volume reflects your organization's established retention policy.

SnapLock volume retention periods

A SnapLock Compliance volume has three retention periods that you can set:

Minimum retention period: The minimum retention period applies to the shortest amount of time the WORM file must be kept in a SnapLock volume. You set this retention period to ensure that applications or users do not assign noncompliant retention periods to retained records in regulatory environments. This option has the following characteristics:

- Existing files that are already in the WORM state are not affected by changes in this volume retention period.
- The minimum retention period takes precedence over the default retention period.
- Until you explicitly reconfigure it, the minimum retention period is 0.

Maximum retention period: The maximum retention period applies to the largest amount of time the WORM file must be kept in a SnapLock volume. You set this retention period to ensure that applications or users do not assign excessive retention periods to retained records in regulatory environments. This option has the following characteristics:

- Existing files that are already in the WORM state are not affected by changes in this volume retention period.
- The maximum retention period takes precedence over the default retention period.
- Until you explicitly reconfigure it, the maximum retention period is 30 years.

Default retention period: The default retention period specifies the retention period assigned to any WORM file on the SnapLock Compliance volume that was not explicitly assigned a retention period. You set this retention period to ensure that a retention period is assigned to all WORM files on the volume, even if users or applications failed to assign a retention period.

Caution

For SnapLock Compliance volumes, the default retention period is equal to the maximum retention period of 30 years. If you do not change either the maximum retention period or the default retention period, for 30 years you will not be able to delete WORM files that received the default retention period.

Setting SnapLock volume retention periods

SnapLock volume retention periods can be specified in days, months, or years. Data ONTAP applies the retention period in a calendar correct method. That is, if a WORM file created on 1 February has a retention period of 1 month, the retention period will expire on 1 March.

Setting the minimum retention period: To set the SnapLock volume minimum retention period, complete the following step.

Step	Action
1	Enter the following command: **`vol options vol_name snaplock_minimum_period period`** *vol_name* is the SnapLock volume name. *period* is the retention period in days (d), months (m), or years (y). **Example:** The following command sets a minimum retention period of 6 months: `vol options wormvol1 snaplock_minimum_period 6m`

Setting the maximum retention period: To set the SnapLock volume maximum retention period, complete the following step.

Step	Action
1	Enter the following command: **`vol options vol_name snaplock_maximum_period period`** *vol_name* is the SnapLock volume name. *period* is the retention period in days (d), months (m), or years (y). **Example:** The following command sets a maximum retention period of 3 years: `vol options wormvol1 snaplock_maximum_period 3y`

Setting the default retention period: To set the SnapLock volume default retention period, complete the following step.

Step	Action
1	Enter the following command: **`vol options vol_name snaplock_default_period [period | min | max]`** *vol_name* is the SnapLock volume name. *period* is the retention period in days (d), months (m), or years (y). `min` is the retention period specified by the `snaplock_minimum_period` option. `max` is the retention period specified by the `snaplock_maximum_period` option. **Example:** The following command sets a default retention period equal to the minimum retention period: `vol options wormvol1 snaplock_default_period min`

Destroying SnapLock volumes and aggregates

When you can destroy SnapLock volumes

SnapLock Compliance volumes constantly track the retention information of all retained WORM files. Data ONTAP does not allow you to destroy any SnapLock volume that contains unexpired WORM content. Data ONTAP does allow you to destroy SnapLock Compliance volumes when all the WORM files have passed their retention dates, that is, expired.

Note

You can destroy SnapLock Enterprise volumes at any time.

When you can destroy SnapLock aggregates

You can destroy SnapLock Compliance aggregates only when they contain no volumes. The volumes contained by a SnapLock Compliance aggregate must be destroyed first.

Destroying SnapLock volumes

To destroy a SnapLock volume, complete the following steps.

Step	Action
1	Ensure that the volume contains no unexpired WORM data.
2	Enter the following command to offline the volume: **`vol offline vol_name`**
3	Enter the following command: **`vol destroy vol_name`**

If there are any unexpired WORM files in the SnapLock Compliance volume, Data ONTAP returns the following message:

```
vol destroy: Volume volname cannot be destroyed because it is a
SnapLock Compliance volume.
```

Destroying SnapLock aggregates

To destroy a SnapLock aggregate, complete the following steps.

Step	Action
1	Using the steps outlined in "Destroying SnapLock volumes" on page 337, destroy all volumes contained by the aggregate you want to destroy.
2	Using the steps outlined in "Destroying an aggregate" on page 172, destroy the aggregate.

Managing WORM data

Transitioning data to WORM state and setting the retention date

After you place a file into a SnapLock volume, you must explicitly commit it to a WORM state before it becomes WORM data. The last accessed timestamp of the file at the time it is committed to WORM state becomes its retention date.

This operation can be done interactively or programmatically. The exact command or program required depends on the file access protocol (CIFS, NFS, etc.) and client operating system you are using. Here is an example of how you would perform these operations using a Unix shell:

Unix shell example: The following commands could be used to commit the document.txt file to WORM state, with a retention date of November 21, 2004, using a Unix shell.

```
touch -a -t 200411210600 document.txt
chmod -w document.txt
```

Note

In order for a file to be committed to WORM state, it must make the transition from writable to read-only in the SnapLock volume. If you place a file that is already read-only into a SnapLock volume, it will not be committed to WORM state.

If you do not set the retention date, the retention date is calculated from the default retention period for the volume that contains the file.

Extending the retention date of a WORM file

You can extend the retention date of an existing WORM file by updating its last accessed timestamp. This operation can be done interactively or programmatically.

Note

The retention date of a WORM file can never be changed to earlier than its current setting.

Determining whether a file is in a WORM state

To determine whether a file is in WORM state, it is not enough to determine whether the file is read-only. This is because to be committed to WORM state, files must transition from writable to read-only *while in the SnapLock volume*.

If you want to determine whether a file is in WORM state, you can attempt to change the last accessed timestamp of the file to a date earlier than its current setting. If the file is in WORM state, this operation fails.

Glossary

ACL Access control list. A list that contains the users' or groups' access rights to each share.

adapter card See host adapter.

aggregate A manageable unit of RAID-protected storage, consisting of one or two plexes, that can contain one traditional volume or multiple FlexVol volumes.

ATM Asynchronous transfer mode. A network technology that combines the features of cell-switching and multiplexing to offer reliable and efficient network services. ATM provides an interface between devices, such as workstations and routers, and the network.

authentication A security step performed by a domain controller for the filer's domain, or by the filer itself, using its /etc/passwd file.

AutoSupport A filer daemon that triggers e-mail messages from the customer site to NetApp, or to another specified e-mail recipient, when there is a potential filer problem.

CIFS Common Internet File System. A file-sharing protocol for networked PCs.

client A computer that shares files on a filer.

cluster A pair of filers connected so that one filer can detect when the other is not working and, if so, can serve the failed filer data. For more information about managing clusters, see the *System Administration Guide*.

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

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

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

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

DAFS Direct Access File System protocol.

degraded mode The operating mode of a filer when a disk is missing from a RAID 4 array, when one or two disks are missing from a RAID-DP array, or when the batteries on the NVRAM card are low.

disk ID number A number assigned by a filer to each disk when it probes the disks at boot time.

disk sanitization A multiple write process for physically obliterating existing data on specified disks in such a manner that the obliterated data is no longer recoverable by known means of data recovery.

disk shelf A shelf that contains disk drives and is attached to a filer.

Ethernet adapter An Ethernet interface card.

expansion card See host adapter.

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

GID Group identification number.

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

host adapter (HA) A SCSI card, an FC-AL card, a network card, a serial adapter card, or a VGA adapter that plugs into a NetApp expansion slot.

hot spare disk A disk installed in the filer that can be used to substitute for a failed disk. Before the disk failure, the hot spare disk is not part of the RAID disk array.

hot swap The process of adding, removing, or replacing a disk while the filer is running.

hot swap adapter An expansion card that makes it possible to add or remove a hard disk with minimal interruption to file system activity.

inode A data structure containing information about files on a filer and in a UNIX file system.

mail host The client host responsible for sending automatic e-mail to NetApp when certain filer events occur.

maintenance mode An option when booting a filer from a system boot disk. Maintenance mode provides special commands for troubleshooting your hardware and your system configuration.

MultiStore An optional software product that enables you to partition the storage and network resources of a single filer so that it appears as multiple filers on the network.

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

NVRAM card An adapter card that contains the filer's NVRAM cache.

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

panic A serious error condition causing the filer to halt. Similar to a software crash in the Windows system environment.

parity disk The disk on which parity information is stored for a RAID 4 disk drive array. In RAID groups using RAID-DP protection, two parity disks store parity and double-parity information. Used to reconstruct data in failed disk blocks or on a failed disk.

PCI Peripheral Component Interconnect. The bus architecture used in newer filer models.

pcnfsd A filer daemon that permits PCs to mount filer file systems. The corresponding PC client software is called (PC)NFS.

qtree A special subdirectory of the root of a volume that acts as a virtual subvolume with special attributes.

RAID Redundant array of independent disks. A technique that protects against disk failure by computing parity information based on the contents of all the disks in an array. NetApp filers use either RAID Level 4, which stores all parity information on a single disk, or RAID-DP, which stores parity information on two disks.

RAID disk scrubbing The process in which a system reads each disk in the RAID group and tries to fix media errors by rewriting the data to another disk area.

SCSI adapter An expansion card that supports SCSI disk drives and tape drives.

SCSI address The full address of a disk, consisting of the disk's SCSI adapter number and the disk's SCSI ID, such as 9a.1.

SCSI ID The number of a disk drive on a SCSI chain (0 to 6).

serial adapter An expansion card for attaching a terminal as the console on some filer models.

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

share A directory or directory structure on the filer that has been made available to network users and can be mapped to a drive letter on a CIFS client.

SID Security identifier.

snapshot An online, read-only copy of an entire file system that protects against accidental deletions or modifications of files without duplicating file contents. Snapshots enable users to restore files and to back up the filer to tape while the filer is in use.

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

trap An asynchronous, unsolicited message sent by an SNMP agent to an SNMP manager indicating that an event has occurred on the filer.

tree quota — A type of disk quota that restricts the disk usage of a directory created by the `quota qtree` command. Different from user and group quotas that restrict disk usage by files with a given UID or GID.

UID — User identification number.

Unicode — A 16-bit character set standard. It was designed and is maintained by the nonprofit consortium Unicode Inc.

vFiler — A virtual filer you create using MultiStore, which enables you to partition the storage and network resources of a single filer so that it appears as multiple filers on the network.

VGA adapter — Expansion card for attaching a VGA terminal as the console.

volume — A file system.

WAFL — Write Anywhere File Layout. The WAFL file system was designed for the NetApp filer to optimize write performance.

WebDAV — Web-based Distributed Authoring and Versioning protocol.

workgroup — A collection of computers running Windows NT or Windows for Workgroups that is grouped for browsing and sharing.

WORM — Write Once Read Many. WORM storage prevents the data it contains from being updated or deleted. For more information about how NetApp provides WORM storage, see “SnapLock Management” on page 327.

Index

Symbols

/etc/messages file 115, 116
/etc/messages, automatic checking of 115
/etc/quotas file
    character coding 291
    Disk field 293
    entries for mapping users 301
    errors in 306
    example entries 290, 298
    file format 289
    Files field 294
    order of entries 290
    quota_perform_user_mapping 302
    quota_target_domain 301
    Soft Disk field 296
    Soft Files field 297
    Target field 292
    Threshold field 295
    Type field 293
/etc/sanitized_disks file 90

A

ACL 341
adapter. *See also* disk adapter and host adapter
aggr commands
    aggr copy 212
    aggr create 158
    aggr offline 164
    aggr online 165
    aggr restrict 165
    aggr status 331
aggregate and volume operations compared 34
aggregate overcommitment 250
aggregates
    adding disks to 34, 168, 170
    aggr0 24
    bringing online 165
    changing states of 35
    changing the RAID type of 122
    changing the size of 34
    copying 35, 165
    creating 28, 36, 158
    creating SnapLock 330
    described 3, 14
    destroying 172
    determining state of 163
    displaying as FlexVol container 37
    displaying disk space of 171
    hot spare disk planning 168
    how to use 14, 154
    maximum limit per appliance 25
    mirrored 4, 155
    new appliance configuration 24
    operations 34
    overcommitting 250
    physically moving between NetApp systems 174
    planning considerations 24
    RAID, changing type 122
    renaming 166
    rules for adding disks to 167
    SnapLock and 330
    states of 162
    taking offline 164
    taking offline, when to 163
    when to put in restricted state 165
assigning disk ownership 45
ATM 341
automatic shutdown conditions 116
Autosupport and SnapLock 329
Autosupport message, about disk failure 116

B

backup
    planning considerations 27
    using qtrees for 258
    with snapshots 8
block checksum disks 2, 44

C

cache hit 233
checksum type 186

block 44, 186
rules 157
zoned 44, 186
CIFS
commands, options cifs.oplocks.enable (enables and disables oplocks) 266
oplocks
changing the settings (options cifs.oplocks.enable) 266
definition of 265
setting for volumes 185, 193
setting in qtrees 258
clones *See* FlexClone volumes
cloning FlexVol volumes 197
commands
disk assign 72
options raid.reconstruct.perf_impact (modifies RAID data reconstruction speed) 131
options raid.reconstruct_speed (modifies RAID data reconstruction speed) 132, 138
options raid.resync.perf_impact (modifies RAID plex resynchronization speed) 133
options raid.scrub.duration (sets duration for disk scrubbing) 138
options raid.scrub.enable (enables and disables disk scrubbing) 138
options raid.verify.perf_impact (modifies RAID mirror verification speed) 134
See also aggr commands, qtree commands, quota commands, RAID commands, storage commands, volume commands
compliance clock
about 332
initializing 332
viewing 333
containing aggregate, displaying 37
continuous media scrub
adjusting maximum time for cycle 144
checking activity 146
description 144
disabling 144, 145
enabling on spare disks 146, 148
converting directories to qtrees 270
converting volumes 33
create_reserved option 252

D

data disks
removing 68
replacing 118
stopping replacement 118
Data ONTAP, upgrading 16, 19, 24, 27, 31, 33
data reconstruction
after disk failure 117
description of 131
data sanitization
planning considerations 25
See also disk sanitization
data storage, configuring 28
degraded mode 69, 116
deleting qtrees 273
destroying
aggregates 38, 172
FlexVol volumes 38
traditional volumes 38
volumes 38, 224
directories, converting to qtrees 270
directory size, setting maximum 39
disk
assign command
modifying 74
use on the FAS270 and 270c appliances 72
commands
aggr status -s (determines number of hot spare disks) 65
df (determines free disk space) 61
df (reports discrepancies) 61
disk scrub (starts and stops disk scrubbing) 136
storage 92
sysconfig -d 58
failures
data reconstruction after 117
predicting 114
RAID reconstruction after 115

without hot spare 116
ownership
about 45
automatically erasing information 77
erasing prior to removing disk 77
modifying assignments 74
software-based 71
undoing accidental conversion to 78
viewing 73
ownership assignment
description 71
modifying 74
sanitization
description 80
licensing 81
limitations 80
log files 90
selective data sanitization 85
starting 82
stopping 85
sanitization, easier on traditional volumes 31
scrubbing
description of 135
enabling and disabling (options raid.scrub.enable) 138
manually running it 139
modifying speed of 132, 138
scheduling 136
setting duration (options raid.scrub.duration) 138
starting/stopping (disk scrub) 136
toggling on and off 138
space, report of discrepancies (df) 61
swap command, cancelling 67
disk speed, overriding 159
disks
adding new to filer 63
adding to a RAID group other than the last RAID group 170
adding to an aggregate 168
adding to NetApp systems 62
adding to system 63
adding to the system 62
addressing 58
assigning 71
assigning ownership of FAS270 systems 45
assigning ownership of of FAS270 and FAS 270c systems 71
available space on new 60
data, removing 68
data, stopping replacement 118
description of 12, 43
determining number of hot spares (sysconfig) 65
failed, removing 66
FC addressing 59
forcibly adding 170
hot spare, adding to NetApp systems 64
hot spare, removing 67
hot spares, displaying number of 65
how initially configured 2, 45
how to use 12
making available 58
ownership of on FAS270 and FAS270c systems 71
portability 27
reasons to remove 66
recognized by Data ONTAP 58
removing 66
replacing
replacing data disks 118
re-using 76
rules for adding disks to an aggregate 167
SCSI addressing 58
software-based ownership 71
speed matching 158
viewing information about 55
when to add 62
double-disk failure
avoiding with media error thresholds 149
RAID-DP protection against 108
without hot spare disk 116
duplicate volume names 212

E

effects of oplocks 265

F

failed disk, removing 66

failure, data reconstruction after disk 117
FAS250 filer, default RAID-4 group size 127
FAS270 filer, assigning disks to 72
FAS270c filer, assigning disks to 72
FC disk addresses 59
Fibre Channel, Multipath I/O 49
file grouping, using qtrees 258
files
    as storage containers 18
    space reservation for 252
files, how used 12
FlexCache volumes
    about 230
    attribute cache timeouts 231
    cache consistency 231
    cache hits and misses 233
    cache objects 230
    creating 238
    description 229
    forward proxy deployment 236
    license requirement 230
    limitations of 233
    reverse proxy deployment 236
    sample deployments 236
    statistics, viewing 242
    volume options 232
    write operation proxy 233
FlexClone volumes
    creating 37, 197
    splitting 201
flexible volumes
    See FlexVol volumes
FlexVol volumes
    about creating 191
    bringing online in an overcommitted aggregate 251
    changing states of 35, 216
    changing the size of 34
    cloning 197
    co-existing with traditional 8
    copying 35
    creating 28, 36, 191
    defined 7
    definition of 178
    described 16
    displaying containing aggregate 203
    how to use 16
    migrating to traditional volumes 205
    operations 190
    resizing 195
    SnapLock and 330
    space guarantees, planning 26
fractional reserve, about 254

G

gFiler gateway LUNs 18
gFiler gateways
    and LUNs 11
    disk ownership 12
    RAID levels supported 3
gFiler gatways
    and LUNs 11
group quotas 276, 281

H

host adapter
    2202 50
    2212 50
    changing state of 94
    storage command 92
    viewing information about 97
hot spare disks
    adding to NetApp systems 64
    displaying number of 65
    removing 67
hub, viewing information about 99

I

inodes 226

L

language
    displaying its code 37
    setting for volumes 39
    specifying the character set for a volume 27
LUNs
    in a SAN environment 18

with gFiler gateways 18
LUNs, how used 11

M

maintenance mode 79, 164
maximum files per volume 226
media error failure thresholds 149
media scrub
adjusting maximum time for cycle 144
continuous 144
continuous. *See also* continuous media scrub
disabling 145
displaying 38
migrating volumes with SnapMover 31
mirror verification, description of 134
mixed security style, description of 261
mode, degraded 69, 116
Multipath I/O
enabling 50
host adapters 50
preventing adapter single-point-of-failure 49
understanding 49

N

naming conventions for volumes 182, 191
NetApp systems
adding disks to 62, 63
adding hot spare disks 64
automatic shutdown conditions 116
determining number of hot spare disks in (sysconfig) 65
running in degraded mode 116
when to add disks 62
NTFS security style, description of 261

O

oplocks
definition of 265
disabling 266
effects when enabled 265
enabling 266
enabling and disabling (options cifs.oplocks.enable) 266
setting for volumes 185, 193
options command, setting filer automatic shutdown 116
overcommitting aggregates 250
overriding disk speed 159

P

parity disks, size of 168
physically transferring data 31
planning
for maximum storage 24
for RAID group size 24
for RAID group type 25
for SyncMirror replication 24
planning considerations 27
backup 27
data sanitization 25
FlexVol space guarantees 26
language 27
qtrees 27
quotas 27
root volume sharing 25
SnapLock volume 25
traditional volumes 27
plex, synchronization 133
plexes
defined 3
described 13
how to use 13
snapshots of 9

Q

qtree commands
qtree create 260
qtree security (changes security style) 263
qtrees
changing security style 263
CIFS oplocks in 257
converting from directories 270
creating 31, 260
definition of 11
deleting 273
described 17, 256
displaying statistics 269

grouping criteria 258
grouping files 258
how to use 11, 17
maximum number 256
planning considerations 27
quotas and changing security style 316
quotas and deleting 316
quotas and renaming 316
reasons for using in backups 258
reasons to create 256
renaming 273
security styles for 261
security styles, changing 263
stats command 269
status, determining 268
understanding 256
qtrees and volumes
changing security style in 263
comparison of 256
security styles available for 261
quota commands
quota logmsg (displays message logging settings) 315
quota logmsg (turns quota message logging on or off) 314
quota off (deactivates quotas) 308
quota off(deactivates quotas) 308
quota off/on (reinitializes quota) 307
quota on (activates quotas) 307
quota on (enables quotas) 307
quota report (displays report for quotas) 326
quota resize (resizes quota) 311
quota reports
contents 320
formats 322
ID and Quota Specifier fields 322
types 319
quota_perform_user_mapping 302
quota_target_domain 301
quotas
/etc/quotas file. *See* /etc/quotas file in the "Symbols" section of this index
/etc/rc file and 279
activating (quota on) 307
applying to multiple IDs 285
canceling initialization 308
changing 309
CIFS requirement for activating 306
conflicting, how resolved 300
console messages 287
deactivating 308
default
advantages of 283
description of 280
examples 299
overriding 280
scenario for use of 280
where applied 280
default UNIX name 305
default Windows name 305
deleting 312
derived 281
disabling (quota off) 308
Disk field 293
displaying report for (quota report) 326
enabling 307
errors in /etc/quotas file 306
example quotas file entries 290, 298
explicit quota examples 298
explicit, description of 277
Files field 294
group 276
group drived from tree 282
group quota rules 290
hard versus soft 277
initialization
canceling 308
description 279
upgrades and 307
message logging
display settings (quota logmsg) 315
turning on or off (quota logmsg) 314
modifying 309
notification when exceeded 287
order of entries in quotas file 290
overriding default 280
planning considerations 27
prerequisite for working 279
qtree
deletion and 316

renaming and 316
security style changes and 316
quota_perform_user_mapping 302
quota_taraget_domain 301
quotas file *See also* /etc/quotas file in the "Symbols" section of this index
reinitializing (quota on) 307
reinitializing versus resizing 309
reports
contents 320
formats 322
types 319
resizing 309, 311
resizing versus reinitializing 309
resolving conflicts 300
root users and 286
SNMP traps when exceeded 287
Soft Disk field 296
Soft Files field 297
soft versus hard 277
Target field 292
targets, description of 276
Threshold field 295
thresholds, description of 277, 295
tree 276
Type field 293
types of reports available, description of 319
types, description of 276
UNIX IDs in 284
UNIX names without Windows mapping 305
user and group, rules for 290
user derived from tree 282
user quota rules 290
Windows
group IDs in 285
IDs in 284
IDs, mapping 301
names without UNIX mapping 305

R

RAID
automatic group creation 108
changing from RAID-4 to RAID-DP 122
changing from RAID-DP to RAID-4 124
changing group size 127
changing RAID type 122
changing the group size option 128
commands
aggr create (specifies RAID group size) 119
aggr status 119
vol volume (changes RAID group size) 122, 128
data reconstruction speed, modifying (options raid.reconstruct.perf_impact) 131
data reconstruction speed, modifying (options raid.reconstruct_speed) 132, 138
data reconstruction, description of 131
description of 105
group size
changing (vol volume) 122, 128
comparison of larger versus smaller groups 112
default size 119
maximum 127
planning 24
specifying at creation (vol create) 119
group size changes
for RAID-4 to RAID-DP 123
for RAID-DP to RAID-4 124
groups
about 13
size, planning considerations 24
types, planning considerations 25
maximum and default group sizes
RAID-4 127
RAID-DP 127
media errors during reconstruction 143
mirror verification speed, modifying (options raid.verify.perf_impact) 134
operations
effects on performance 130
types you can control 130
options
setting for aggregates 40
setting for traditional volumes 40
parity checksums 2
plex resynchronization speed, modifying (options raid.resync.perf_impact) 133

reconstruction
media error encountered during 142
reconstruction of disk failure 115
status displayed 150
throttling data reconstruction 131
type
changing 122
descriptions of 106
verifying 126
verifying RAID type 126
verifying the group size option 129
RAID groups
adding disks 170
RAID-4
maximum and default group sizes 127
See also RAID
RAID-DP
maximum and default group sizes 127
See also RAID
RAID-level scrub performing
on aggregates 39
on traditional volumes 39
rapid RAID recovery 114
reallocation, running after adding disks for LUNs 171
reconstruction after disk failure, data 117
reliability, improving with MultiPath I/O 49
renaming
aggregates 39
flexible volumes 39
traditional volumes 39
volumes 39
renaming qtrees 273
resizing FlexVol volumes 195
restoring
with snapshots 8
restoring data with snapshots 256
restoring data, using qtrees for 258
root volume, setting 40
rooted directory 270

S

SCSI disk addresses 58
security styles
changing of, for volumes and qtrees 259, 263
for volumes and qtrees 261
mixed 261
NTFS 261
setting for volumes 185, 193
types available for qtrees and volumes 261
UNIX 261
shutdown conditions 116
single 149
single-disk failure
without hot spare disk 107, 116
SnapLock
about 328
aggregates and 330
Autosupport and 329
compliance clock
about 332
initializing 332
viewing 333
creating aggregates 330
creating traditional volumes 330
data, moving to WORM state 339
destroying aggregates 338
destroying volumes 337
files, determining if in WORM state 340
FlexVol volumes and 330
how it works 328
licensing 329
replication and 329
retention dates
extending 339
setting 339
retention periods
default 334
maximum 334
minimum 334
setting 335
when to set 334
volume retention periods *See* SnapLock retention periods
volumes
creating 37
planning considerations 25
when you can destroy aggregates 337
when you can destroy volumes 337

WORM requirements 332
write_verify option 331
SnapLock Compliance, about 328
SnapLock Enterprise, about 328
SnapMirror software 9
SnapMover
described 55
volume migration, easier with traditional volumes 31
snapshot 8
software-based disk ownership 71
space guarantees
about 247
changing 250
setting at volume creation time 249
space management
about 244
how to use 245
traditional volumes and 248
space reservations
about 252
enabling for a file 253
querying 253
speed matching of disks 158
splitting FlexClone volumes 201
status
displaying aggregate 38
displaying FlexVol 38
displaying traditional volume 38
storage commands
changing state of host adapter 94
disable 94, 95
enable 94, 95
managing host adapters 92
reset tape drive statistics 103
show 96
syntax 92
viewing information about
disks 55
host adapters 97
hubs 99
media changers 101
primary and secondary paths 56
storage subsystem 96
supported tape drives 102
switch ports 101
switches 101
tape drives 102
storage, maximizing 24
swap disk command
cancelling 67
SyncMirror replica, creating 37
SyncMirror, planning for 24

T

thin provisioning. *See* aggregate overcommitment
traditional volumes
adding disks 34
changing states of 35, 216
changing the size of 34
copying 35
creating 31, 36, 182
creating SnapLock 330
definition of 16, 178
how to use 16
migrating to FlexVol volumes 205
operations 181
planning considerations, transporting disks 27
reasons to use 31
See also volumes
space management and 248
transporting 27
transporting between NetApp systems 187
upgrading to Data ONTAP 7.0 27
transporting disks, planning considerations 27
tree quotas 276

U

UNICODE options, setting 40
UNIX security style, description of 261
uptime, improving with MultiPath I/O 49

V

volume and aggregate operations compared 34
volume commands
maxfiles (displays or increases number of files) 227, 249, 253
vol create (creates a volume) 160, 183, 191

vol create (specifies RAID group size) 119
vol destroy (destroys an off-line volume) 195, 199, 201, 203, 224
vol lang (changes volume language) 215
vol offline (takes a volume offline) 220
vol online (brings volume online) 221
vol rename (renames a volume) 223
vol restrict (puts volume in restricted state) 165, 221
vol status (displays volume language) 214
vol volume (changes RAID group size) 128
volume names, duplicate 212
volume operations 34, 179, 204
volume-level options, configuring 40
volumes
aggregates as storage for 6
as a data container 5
attributes 10, 25
bringing online 165, 221
bringing online in an overcommitted aggregate 251
cloning FlexVol 197
common attributes 15
conventions of 157
converting from one type to another 33
creating (vol create) 157, 160, 183, 191
creating FlexVol volumes 191
creating traditional 182
creating traditional SnapLock 330
destroying (vol destroy) 195, 199, 201, 203, 224
destroying, reasons for 195, 224
displaying containing aggregate 203
duplicate volume names 212
flexible. *See* FlexVol volumes
how to use 15
increasing number of files (maxfiles) 227, 249, 253
language
changing (vol lang) 215
choosing of 213
displaying of (vol status) 214
planning 27
limits on number 179
maximum limit per appliance 25
maximum number of files 226
migrating between traditional and FlexVol 205
mirroring of, with SnapMirror 9
naming conventions 182, 191
number of files, displaying (maxfiles) 227
operations for FlexVol 190
operations for traditional 181
operations, general 204
post-creation changes 185, 193
renaming 223
renaming a volume (vol rename) 166
resizing FlexVol 195
restricting 221
root, planning considerations 25
root, setting 40
security style 185, 193
SnapLock, creating 37
SnapLock, planning considerations 25
specifying RAID group size (vol create) 119
taking offline (vol offline) 220
traditional. *See* traditional volumes
volume state, definition of 216
volume state, determining 219
volume status, definition of 216
volume status, determining 219
when to put in restricted state 220
volumes and qtrees
changing security style 263
comparison of 256
security styles available 261
volumes, traditional
co-existing with FlexVol volumes 8

W

WORM
data 328
determining if file is 340
requirements 332
transitioning data to 339

Z

zoned checksum disks 2, 44
zoned checksums 186