

Data ONTAP® 7.1

Upgrade 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 210-01083_B0
December 2005

Copyright and trademark information

Copyright information

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

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

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

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

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

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

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

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

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

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

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

This product includes software developed by the University of California, Berkeley and its contributors.

4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

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

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

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

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

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

Portions Copyright © 2001, Sitraka Inc.

Portions Copyright © 2001, iAnywhere Solutions

Portions Copyright © 2001, i-net software GmbH

Portions 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, Center-to-Edge, ComplianceClock, ComplianceJournal, ContentDirector, 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, SpinServer, and SpinStor are registered trademarks of Spinnaker Networks, LLC in the United States and/or other countries. SpinAV, SpinManager, SpinMirror, SpinRestore, and SpinShot 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	vii
Chapter 1	Overview of the Upgrade Process	1
Chapter 2	Before You Upgrade: What Is Required.	3
	System requirements	4
	Upgrade host requirements	6
	Planning for upgrading clusters.	8
	Planning for upgrading storage systems running SnapMirror	12
	Resolving upgrade issues	16
	Upgrading to Data ONTAP 7.1 or later	18
	Upgrading from releases earlier than Data ONTAP 7.0	24
	Upgrading from releases earlier than Data ONTAP 6.5	34
	Upgrading from releases earlier than Data ONTAP 6.4	46
	Upgrading from releases earlier than Data ONTAP 6.2	47
	Upgrading from releases earlier than Data ONTAP 6.1	48
	Upgrading iSCSI hosts	50
Chapter 3	Upgrading NetApp Firmware	59
	Upgrading the system firmware	60
	Upgrading disk firmware	67
Chapter 4	Upgrading Data ONTAP Software.	73
	Prerequisites to software upgrade	74
	Installing Data ONTAP system files	77
	UNIX client steps	78
	Windows client steps.	82
	Using the Upgrade Utility	86
	HTTP server steps	87
	Downloading Data ONTAP to single systems and rebooting	92
	Downloading Data ONTAP to clustered systems and rebooting.	98
	Choosing a clustered upgrade procedure.	99

	Nondisruptive upgrades from the Data ONTAP 7.0 release family . . .	101
	Nondisruptive upgrades from the Data ONTAP 6.5 release family . .	109
	Nondisruptive upgrades within release families	117
	Standard upgrades	124
	Concluding the upgrade	135
	Where to go from here	137
Chapter 5	Reverting Storage System Software	139
	Reverting to a previous release	140
	Reversion issues between release families	141
	Reverting to the Data ONTAP 7.0 release family	143
	Reverting to the Data ONTAP 6.5 release family	145
	Reverting to the Data ONTAP 6.4 release family	152
	Reverting to Data ONTAP 6.3 or earlier.	157
Appendix A	Creating Diskettes	159
	Creating diskettes from NOW	160
	Creating diskettes from the CD-ROM	164
Appendix B	DAVE Macintosh Client resource.frk Conversion	167
Appendix C	Understanding Release Family Upgrades	171
Appendix D	Optimizing Service Availability During Upgrades	173
	Index	177

Preface

Introduction

This guide describes how to upgrade NetApp® storage systems to the latest release of Data ONTAP® firmware and software. It covers all storage system models that run Data ONTAP 7.1 and higher; see the *Data ONTAP 7.1 Release Notes* for more information about supported storage systems.

Note

Upgrade procedures for gateway filer (V-Series, formerly known as gFiler™) systems are covered in V-Series documentation. The information in this guide does not apply to V-Series systems.

Audience

This guide is for system administrators who are familiar with operating systems that run on the storage system's clients, such as UNIX®, Windows NT®, Windows® 2000, and Windows XP. It assumes that you know how to use a Web browser.

Accessing Data ONTAP man pages

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

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

Man pages can be viewed in the following ways:

- ◆ At the storage system command line, by entering
`man command_or_file_name`
- ◆ From the FilerView® main navigational page
- ◆ In *Commands: Manual Page Reference*, Volumes 1 and 2

Note

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

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

Terminology

Storage systems that run Data ONTAP are sometimes referred to as filers, storage appliances, appliances, or systems. The name of the graphical user interface for Data ONTAP, FilerView, reflects one of these common usages.

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, or clicking in a field in a graphical interface and typing information into it.

FilerView as an alternative to commands

Tasks you perform as a Data ONTAP administrator can be performed by entering commands at the storage system console, in configuration files, or through a Telnet session or a Remote Shell connection.

Another method of performing many common tasks is to use the FilerView® graphical management interface for viewing and managing storage system from a Web browser. FilerView comes with every storage system, is easy to use, and includes Help that explains Data ONTAP features and how to work with them in FilerView.

For more information about accessing a storage system with FilerView, and about FilerView Help, see the *System Administration Guide*.

Command conventions

In examples that illustrate commands executed on a UNIX workstation, 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
<i>Italic font</i>	Words or characters that require special attention. Placeholders for information you must supply. For example, if the guide says to enter the <code>arp -d hostname</code> command, you enter the characters “arp -d” followed by the actual name of the host. Book titles in cross-references.
Monospaced font	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 two types of special messages:

Note

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

Attention

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

Overview of the Upgrade Process

1

The upgrade process

The upgrade process requires the following main steps.

Step	Task
1	Plan your upgrade by familiarizing yourself with requirements and issues described in Chapter 2, “Before You Upgrade: What Is Required,” on page 3. Plan to do the following: ◆ Understand any requirements for upgrading to Data ONTAP 7.1 from your existing software. ◆ Note any potential changes to your system after the upgrade. ◆ If you have clustered storage systems, select an upgrade method. ◆ If you run the SnapMirror® software, identify storage systems with destination and source volumes.
2	If necessary, perform any required preliminary upgrades before upgrading to Data ONTAP 7.1. Required upgrades might include the following: ◆ Performing an intermediate upgrade, as described in “Resolving upgrade issues” on page 16 ◆ “Upgrading the system firmware” on page 60 ◆ “Upgrading disk firmware” on page 67
3	Obtain the Data ONTAP system files from the NetApp® on the Web™ (NOW) site at http://now.netapp.com/ or the Data ONTAP CD-ROM.
4	Install the 7.1 system files and download them to your storage systems as described in Chapter 4, “Upgrading Data ONTAP Software,” on page 73.
5	Conclude the upgrade with any of the tasks discussed in “Concluding the upgrade” on page 135.

Reverting to a previous release

For information about reverting to a previous Data ONTAP release, see Chapter 5, “[Reverting Storage System Software](#),” on page 139.

About this chapter

This chapter contains information to help you ensure that your storage systems meet the requirements for upgrading to the new version of Data ONTAP. Additionally, this chapter provides information about some issues you might encounter during an upgrade. Because new features are introduced in each release of Data ONTAP, you might encounter issues when upgrading, especially if you are not upgrading from the immediately previous version of Data ONTAP.

Information about bug fixes

For a list of bugs fixed in this release, access Bugs Online by completing the following steps.

Step	Action
1	Use a Web browser to go to http://now.netapp.com/cgi-bin/bol/ .
2	Click Service & Support.
3	Click Access Software Bug Tools.

D or P patches might not be included: Data ONTAP 7.1 might not contain all the bug fixes in preceding D or P patch releases.

Attention

If your storage system is running a patch release, compare the bug fixes for this release to the bug fixes for your patch release to determine whether you should upgrade. For information about these bug fixes, see Bugs Online at <http://now.netapp.com/cgi-bin/bol/>.

Topics in this chapter

This chapter describes the following topics:

- ◆ “[System requirements](#)” on page 4
- ◆ “[Upgrade host requirements](#)” on page 6
- ◆ “[Planning for upgrading clusters](#)” on page 8
- ◆ “[Planning for upgrading storage systems running SnapMirror](#)” on page 12
- ◆ “[Resolving upgrade issues](#)” on page 16
- ◆ “[Upgrading iSCSI hosts](#)” on page 50

System requirements

Data ONTAP versions supported

To determine whether your storage system supports Data ONTAP 7.1, complete the following steps.

Step	Action
1	Use a Web browser to go to http://now.netapp.com/ .
2	Click Service & Support.
3	Click Download Software.
4	In the Select Platform list box in the Data ONTAP product row, select your storage system type.
5	Click Go. Result: You see a list of the versions of Data ONTAP supported by your storage system platform. If Data ONTAP 7.1 is listed as a release for your storage system, you can upgrade to it.

Storage system firmware availability

To view system firmware versions available for your storage systems, complete the following steps.

Note

F87 storage systems use firmware that does not need to be upgraded.

Step	Action
1	Use a Web browser to go to http://now.netapp.com/ .
2	Click Service & Support.
3	Click Download Software.
4	Scroll down to Firmware & MIBs.

Step	Action	
5	If you are upgrading ...	Then...
	A CompactFlash®-based storage system	Click System Firmware/Diagnostics for CompactFlash-based Filers, gFilers, & NearStores.
	A diskette-based storage system	Click System Firmware for Filers & NetCache.

For more information about system firmware updates, see “[Upgrading the system firmware](#)” on page 60.

Fibre Channel disk firmware requirement

To view the latest required firmware revisions for Fibre Channel disk drives, complete the following steps.

Step	Action
1	Use a Web browser to go to http://now.netapp.com/ .
2	Click Service & Support.
3	Click Disk Support Matrix & Firmware. Result: The Disk Drive & Firmware matrix appears, where you can click the link to the firmware revision needed for your drives.

For more information about disk firmware requirements and updates, see “[Upgrading disk firmware](#)” on page 67.

Upgrade host requirements

About upgrade hosts

An upgrade host is a system from which you install the Data ONTAP upgrade package on your storage system. There are several types of upgrade host, which determine how you install an upgrade package:

- ◆ From a Windows or UNIX client using the CIFS or NFS protocols
You copy the Data ONTAP upgrade package to a client directory that is mapped as a drive (Windows) or mounted (UNIX) from the storage system being upgraded. You then use commands on the client system to unpack the upgrade package.

Note

To use this method, the CIFS or NFS protocols must be licensed on the storage system.

The iSCSI protocol provides limited CIFS functionality. However, it is not sufficient to create shares on the storage system for the upgrade package, or to use Windows client commands to proceed with the upgrade process.

- ◆ From a Windows or UNIX client using the Upgrade Utility
You use the Data ONTAP Upgrade Utility to copy the Data ONTAP upgrade package to a storage system and to unpack the upgrade package. This method can be used if you do not have CIFS or NFS licensed on your storage system; for example, if you are serving data using the FCP or iSCSI protocol.

Note

You must be running Java on the storage system to use the Upgrade Utility.

- ◆ From an HTTP server
You copy the Data ONTAP upgrade package to an HTTP server, then copy it to the storage system and unpack it using commands on the storage system console. This method can be used if you do not have CIFS or NFS licensed on your storage system, if you want to distribute Data ONTAP upgrade packages to multiple storage systems using an HTTP server, or if you want to use installation scripts. For information about the console, see the *System Administration Guide*.

The host from which you upgrade your storage system must have access to at least one of the following:

- ◆ The NOW site
- ◆ A CD-ROM drive
- ◆ An HTTP server containing Data ONTAP updates

For detailed instructions about preparing an upgrade host, see “[Installing Data ONTAP system files](#)” on page 77.

Requirements for the client

You must be able to administer the storage system from the client.

The following lists the requirements for each type of client.

- ◆ UNIX client—any available version of UNIX
- ◆ Windows client—any version of Windows (Windows NT must be version 3.51 or later.)

The client is usually the storage system’s administration (admin) host. For information about admin hosts, see the *System Administration Guide*.

Requirements for the HTTP server

You must be able to serve the upgrade package from the HTTP server, and you must know the exact URL (including any necessary host and port information) to enter at the storage system console.

Planning for upgrading clusters

Before upgrading clustered storage systems

Before upgrading your clustered storage systems, verify that the cluster is properly configured by using the Cluster Configuration Checker. Identify and resolve any cluster configuration issues before continuing with the upgrade.

The Cluster Configuration Checker is a Perl script that detects errors in the configuration of a pair of clustered NetApp storage systems. The script is available on the NetApp on the Web™ (NOW) site at http://now.netapp.com/NOW/download/tools/cf_config_check/.

See also the chapter on cluster management in the *Cluster Installation and Administration Guide*.

Selecting an upgrade method for your cluster

The following methods are available for upgrading cluster partners:

- ◆ Nondisruptive (CompactFlash-based systems only)
- ◆ Standard

Depending on the nature of the upgrade, nondisruptive cluster upgrades ensure better storage system availability than single system or standard cluster upgrades.

When planning for your upgrade, evaluate the advantages and disadvantages of each method. Each method is described in this section.

Note

You cannot use the nondisruptive method to upgrade clusters to Data ONTAP 7.1 or later from a Data ONTAP release earlier than 6.5.3. To do this upgrade for releases earlier than 6.5.3, you must use the standard method of upgrading.

About nondisruptive cluster upgrades

In a nondisruptive upgrade (also called a *rolling upgrade*), clustering technology is used to enable a takeover storage system to assume the functions of the “failed” partner while it is being upgraded. There is a takeover and giveback operation for each cluster node. Because the cluster partner node fulfills service requests during the “failed” storage system’s upgrade, no disruption in service is experienced by the clients.

In addition, because the takeover storage system assures continuous availability of the “failed” system’s disks, more extensive upgrades requiring a system halt—such as system firmware upgrades and hardware adapter replacements—can be performed without disrupting services based on stateless protocols.

Note

For storage systems with FCP or iSCSI hosts, nondisruptive upgrades within Data ONTAP release families (for example, from Data ONTAP 7.1 to Data ONTAP 7.1.1) are currently supported and can be performed as documented in this guide. Nondisruptive upgrades between Data ONTAP release families (for example, from Data ONTAP 7.0 to Data ONTAP 7.1) will be supported in the future through NetApp Professional Services to ensure that the upgrade is completed seamlessly for the FCP or iSCSI hosts.

Nondisruptive clustered upgrades are recommended to minimize service disruption when one or more of the following is being performed:

- ◆ Upgrades to Data ONTAP 7.1 from the previous release family (for example, from 6.5.3 to 7.1)
- ◆ Data ONTAP upgrades within a release family (for example, from 7.1 to 7.1.1)
- ◆ System firmware upgrades
- ◆ Certain hardware upgrades

Note

See the *Cluster Installation and Administration Guide* for more information about changing system hardware nondisruptively.

Before undertaking a nondisruptive upgrade procedure, make sure that you do the following:

- ◆ Familiarize yourself with the nondisruptive upgrade method.
Nondisruptive upgrades require more planning and procedural steps than standard upgrades. Make sure to thoroughly review this document before starting a nondisruptive upgrade. See especially Appendix D, “[Optimizing Service Availability During Upgrades](#),” on page 173.
- ◆ Ensure that your cluster is optimally configured and functioning correctly.
You can verify that your clustered storage systems are properly configured by running the Cluster Configuration Checker as described in “[Before upgrading clustered storage systems](#)” on page 8.
- ◆ Ensure that your clients are optimally configured and functioning correctly.
Configure client timeout settings to optimize availability as described in “[Service and protocol considerations](#)” on page 174.

- ◆ Disable automatic cluster giveback
If you have enabled automatic giveback by setting the `cf.giveback.auto.enable` option to On, you must set this option to Off before upgrading your cluster using the nondisruptive method. After the upgrade procedure, reset this option to On (if desired).
- ◆ If you need disk firmware upgrades in addition to the Data ONTAP upgrade, ensure that all disks on your system are in RAID-DP™ or mirrored RAID4 aggregates.
Disk firmware upgrades take place automatically in the background when RAID-DP protection is configured. Services and data continue to be available during the disk firmware upgrade. For more information, see [“Service availability during disk firmware upgrades”](#) on page 67.

You should not use nondisruptive upgrades in the following circumstances:

- ◆ You have storage systems actively serving CIFS to clients.
Because CIFS is session-oriented, sessions must be terminated before upgrade procedures to prevent data loss.
- ◆ You have storage systems actively serving File Transfer Protocol (FTP) or Network Data Management Protocol (NDMP) clients that cannot be postponed.
Because these protocols are session-oriented, outstanding sessions must complete, and these services must be disabled in order to use nondisruptive upgrades.
- ◆ You need to upgrade disk firmware and you have RAID4 aggregates on your system.
Standard disk firmware upgrades automatically take disks in RAID4 aggregates offline until the upgrade is complete. Services and data are unavailable until they are back online. For more information, see [“Service availability during disk firmware upgrades”](#) on page 67.

Note

If you upgrade RAID protection to RAID-DP, disk firmware upgrades take place in the background and are non-disruptive.

- ◆ You use diskette-based clustered storage systems (F810, F820, and F840 storage systems).
On diskette-based clustered storage systems, the takeover and giveback operations are much slower than on CompactFlash-based systems and possibly slower than rebooting.

For these conditions, standard upgrades are recommended.

About standard cluster upgrades

In a standard upgrade, the cluster is disabled and each node is updated. When the cluster is disabled, each node behaves as a single-node (non-clustered) storage system; in other words, system services associated with that node are interrupted for as long as it takes the system to reboot.

In general, standard cluster upgrades are most appropriate if you schedule regular downtime for your storage systems. Other maintenance upgrades, such as system firmware and hardware, can also take place when the cluster is disabled.

Standard cluster upgrades must be used for upgrades to Data ONTAP 7.1 or later from releases earlier than Data ONTAP 6.5.3.

Planning for upgrading storage systems running SnapMirror

Upgrade requirements for SnapMirror

If you are upgrading Data ONTAP on storage systems that are running the SnapMirror® software, you must upgrade the systems that have SnapMirror destination volumes *before* you upgrade the systems that have SnapMirror source volumes. For SnapMirror volume replication, the destination volume must run under a version of Data ONTAP equal to or later than that of the SnapMirror source volume.

Note

This requirement applies to both asynchronous and synchronous SnapMirror for volume replication.

The requirement does not apply to SnapMirror for qtree replication, SnapVault, or data restoration from tape using the `restore` command. However, make sure that when you upgrade systems that use these features, you should upgrade your SnapMirror destination systems, SnapVault® secondary systems and restoration target systems before the corresponding source systems to maintain backward compatibility.

To upgrade asynchronous or synchronous SnapMirror for volume replication, use the following process.

Stage	Description
1	Identify any destination volumes.
2	Upgrade the storage systems with destination volumes, beginning with the furthest one in the topology (that is, the last system in a series of cascading destinations).
3	Upgrade the storage system with the source volume.

The following sections provide more detail about:

- ◆ [“Identifying SnapMirror destination volumes”](#) on page 13
- ◆ [“Upgrading storage systems running synchronous SnapMirror”](#) on page 14
- ◆ [“Why you must plan for SnapMirror upgrades”](#) on page 14

- ◆ [“Upgrading storage systems that are mirroring volumes to each other”](#) on page 15
- ◆ [“For more information”](#) on page 15

Attention

You must upgrade the storage systems that have SnapMirror destination volumes *before* upgrading those that have source volumes. If you upgrade the source volumes first, SnapMirror volume replication is disabled. To re-enable SnapMirror volume replication, you must downgrade the source system or upgrade the destination system, so that the version of Data ONTAP on the source system is earlier or the same as that on the destination system.

**Identifying
SnapMirror
destination
volumes**

SnapMirror source volumes can be replicated to single or multiple destination volumes. (Replication to multiple destination volumes is also referred to as replication to *cascading destinations*.) When you upgrade Data ONTAP, you must identify all destination volumes and then upgrade the storage systems on which they reside before upgrading the system where the source volume resides. In addition, when you upgrade storage systems in a cascading series, you should upgrade the systems in order, beginning with the destination systems furthest in your topology from the source system.

Note

Upgrading to Data ONTAP 7.0 or later preserves volumes created with earlier release families as traditional volumes. With the Data ONTAP 7.0 release family, you can also create FlexVol™ volumes. When running SnapMirror software, the source and destination volumes must be of the same type, either traditional or flexible. For information about traditional and FlexVol volumes, see the *Storage Management Guide* and [“About volume types after upgrading to Data ONTAP 7.0 or later”](#) on page 28 in this guide.

To identify storage systems with SnapMirror destination volumes, complete the following step.

Step	Action
1	On the storage system with the source volume, enter the following command: <pre>snapmirror destinations</pre> Result: The <code>snapmirror</code> command lists all destination volumes, including cascaded destinations. For more information about the <code>snapmirror</code> command and options, see the <code>snapmirror(1)</code> man page.

Upgrading storage systems running synchronous SnapMirror

Use the same process to upgrade storage systems running synchronous SnapMirror as described in [“Upgrade requirements for SnapMirror”](#) on page 12.

However, synchronous SnapMirror requires that the source and destination run the same version of Data ONTAP. Hence, when you upgrade a destination storage system in a synchronous mirror, SnapMirror goes into asynchronous mode. When in asynchronous mode, the source system asynchronously replicates data to the destination every minute until a synchronous replication can be reestablished; that is, until the source system is upgraded so that the same Data ONTAP version is running on destination and source systems.

Why you must plan for SnapMirror upgrades

When you upgrade Data ONTAP on storage systems that have a mirrored relationship with each other using SnapMirror, the order in which you upgrade the storage systems is critical. If you do not upgrade in the correct order, you can cause a lapse in SnapMirror replication coverage.

A SnapMirror transfer is initiated when the destination storage system reads a Snapshot™ copy of the source system. Therefore, the destination system must be upgraded first, so that it can read the snapshots of the earlier release. If the source system is upgraded first, the destination system cannot read the source system’s file system, thereby disabling SnapMirror.

Upgrading storage systems that are mirroring volumes to each other

SnapMirror can be configured to enable two storage systems to mirror each other's volumes. In this case, each storage system is both a source and a destination system. For example, System A can mirror volumes to System B, and System B can mirror volumes to System A.

In this configuration, there is logically no way to update both destinations before the corresponding source systems. Therefore, to upgrade Data ONTAP on storage systems that are mirroring volumes to each other, you must disable the mirror, upgrade each system, and reenabling the mirror.

For more information

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

Resolving upgrade issues

What you need to know

Every Data ONTAP release family has unique requirements concerning upgrades. Before you decide whether to upgrade, you should understand

- ◆ Issues you must resolve before upgrading to the new release
- ◆ New system behavior after upgrading to the new release

Because significant new features are introduced in each new Data ONTAP release family, you might encounter issues when upgrading to a new release family, especially if you are not upgrading from the immediately previous version of Data ONTAP.

For example, if you are upgrading from a release in the 6.4 family to Data ONTAP 7.1, you must review and resolve upgrade issues associated with the 6.5 and 7.0 release families (but not 6.4 or earlier) before upgrading to 7.1.

For more information, see Appendix C, “[Understanding Release Family Upgrades](#),” on page 171.

The following sections discuss upgrade issues and new system behavior associated with release families:

- ◆ “[Upgrading to Data ONTAP 7.1 or later](#)” on page 18
- ◆ “[Upgrading from releases earlier than Data ONTAP 7.0](#)” on page 24
- ◆ “[Upgrading from releases earlier than Data ONTAP 6.5](#)” on page 34
- ◆ “[Upgrading from releases earlier than Data ONTAP 6.4](#)” on page 46
- ◆ “[Upgrading from releases earlier than Data ONTAP 6.2](#)” on page 47
- ◆ “[Upgrading from releases earlier than Data ONTAP 6.1](#)” on page 48

Required intermediate upgrades

Depending on the version of Data ONTAP software you currently have installed and the model of your storage system, you might have to upgrade to an intermediate release before upgrading to Data ONTAP 7.1.

- ◆ If your CompactFlash-based storage system is running Data ONTAP 6.4.1 or earlier:
You must upgrade to the latest release in the 6.4 family before upgrading to 7.1.

- ◆ If your diskette-based storage system is running Data ONTAP 6.2 or earlier:
You must upgrade to the latest release in the 6.3 family before upgrading to 7.1.

For more information, see “[Upgrading from releases earlier than Data ONTAP 6.5](#)” on page 34.

Upgrading to Data ONTAP 7.1 or later

Issues you must resolve before upgrading

The following sections discuss upgrade issues that you must resolve before upgrading to the Data ONTAP 7.1 release family:

- ◆ [“SnapVault relationships to WORM volumes cannot be updated after upgrading”](#) on page 18
- ◆ [“iSCSI targets unavailable after upgrading to Data ONTAP 7.1”](#) on page 18
- ◆ [“DAFS services no longer supported”](#) on page 19
- ◆ [“SnapLock Compliance license update”](#) on page 20
- ◆ [“FlexClone license update”](#) on page 20
- ◆ [“Space guarantees now enabled on FlexClone volumes”](#) on page 21
- ◆ [“SnapDrive for Windows upgrade to version 3.1.1”](#) on page 21
- ◆ [“Change in CIFS SMB signing default behavior”](#) on page 22
- ◆ [“Data ONTAP 7.1 upgrades could take longer than usual”](#) on page 22
- ◆ [“Recommendations for all storage systems upgrading to Data ONTAP 7.1”](#) on page 22

SnapVault relationships to WORM volumes cannot be updated after upgrading

Data ONTAP 7.1 includes new SnapVault functionality that provides better support for SnapVault relationships to WORM volumes. If you created SnapVault relationships to WORM volumes in previous Data ONTAP releases, the relationships continue to exist after the upgrade. However, because of changes in SnapVault technology, they cannot be updated.

You are advised to identify existing SnapVault relationships to WORM volumes, destroy them before upgrading, then reconfigure these SnapVault relationships after the upgrade has completed. For more information about establishing SnapVault relationships to SnapLock™ volumes, see the *Data Protection Online Backup and Recovery Guide*.

iSCSI targets unavailable after upgrading to Data ONTAP 7.1

In Data ONTAP 7.1, the iSCSI target portal group tags have changed from the values used in previous releases. You must reconfigure any iSCSI hosts that use these tags or they will not be able to identify the iSCSI target provided by the NetApp storage system.

Attention

HP-UX and Linux hosts must be reconfigured *before* upgrading to Data ONTAP 7.1.

For information about reconfiguring iSCSI hosts to identify Data ONTAP 7.1 iSCSI targets, see “[Upgrading iSCSI hosts](#)” on page 50.

DAFS services no longer supported

Beginning with Data ONTAP 7.1, Direct Access File System (DAFS) services are no longer supported.

To ensure that data previously exported with DAFS services remains available after upgrading to Data ONTAP 7.1, you are advised to transition critical data services from DAFS to other protocols *before* upgrading. To do so, complete the following steps.

Step	Action
1	Identify volumes and qtrees exported with DAFS by entering the following command: dafs export list
2	For each entity identified with the <code>dafs export</code> command, reassign the export to another service (such as CIFS or NFS). For current information about data access services, see the <i>File Access and Protocol Management Guide</i> .
3	Identify VI NIC interfaces that use the DAFS service by entering the following command: dafs status -n
4	Halt the DAFS service by entering the following command: dafs stop
5	For each VI NIC interface identified with the <code>dafs status</code> command, detach the interface from the DAFS service by entering the following command: dafs detach <i>interface_name</i>
6	Make appropriate changes on DAFS client systems to ensure continuity of data access.

For more information about configuring and managing DAFS, see the *File Access Management Guide* and dafs man pages for the Data ONTAP release on which you originally configured DAFS services.

SnapLock Compliance license update

If you have a SnapLock™ Compliance license and you are upgrading to Data ONTAP 7.1 from a Data ONTAP version earlier than 7.0.0.1, you must obtain a unique, updated SnapLock Compliance license code from the Protocol Licenses page, located on the NOW site at the following URL:

<https://now.netapp.com/eservice/agree.do?moduleName=PROTOCOL>

Previous SnapLock Compliance license codes have been replaced with updated license codes on the Protocol Licenses page.

Note

There are two types of SnapLock licenses: SnapLock Compliance and SnapLock Enterprise. When you display licenses for your storage system, the Snaplock Compliance license is listed as “snaplock” and the Snaplock Enterprise license as “snaplock_enterprise”. This update concerns SnapLock Compliance licenses only.

If a storage system is upgraded to Data ONTAP 7.0.0.1 or later without changing the license code, the attributes and SnapLock protection on existing SnapLock files (and SnapLock volumes) are not affected in any way. However, since the old license is no longer functional, you cannot commit any new files to SnapLock protected status (WORM—Write-Once-Read-Many) until the license code is updated. Also, you cannot create new WORM volumes until the license code is updated.

For more information about displaying and managing licenses, see the *System Administration Guide*.

FlexClone license update

If you created FlexClone™ volumes (clones of FlexVol volumes) in a Data ONTAP 7.0.x release and you are upgrading to Data ONTAP 7.1, you must obtain a unique FlexClone license code from the Protocol Licenses page, located on the NOW site at the following URL:

<https://now.netapp.com/eservice/agree.do?moduleName=PROTOCOL>

The FlexClone license was not enforced in Data ONTAP 7.0.0x releases (you could create FlexClone volumes without a license). However, the license is enforced in Data ONTAP 7.0.1 and later releases. When you upgrade to Data

ONTAP 7.1, you do not need a license to access FlexClone volumes created in previous releases, but you cannot create new FlexClone volumes without a license.

For more information about FlexClone volumes, see the *Storage Management Guide*.

Space guarantees now enabled on FlexClone volumes

If you created FlexClone volumes in releases of Data ONTAP earlier than 7.1 and you specified space guarantees in the parent volumes, these guarantees were inherited but not enabled in the FlexClone volumes. FlexClone volumes in earlier releases therefore were not subject to aggregate overcommitment requirements.

Beginning with Data ONTAP 7.1, space guarantees are enabled in FlexClone volumes, including FlexClone volumes that were created in earlier releases. To ensure that existing FlexClone volumes can obtain guaranteed space when mounted after upgrading, you are advised to review the volume guarantees of any existing FlexClone volumes *before* upgrading to Data ONTAP 7.1.

You can change space guarantees to enable or disable aggregate overcommitment using the `vol options guarantee` command. For more information about FlexClone volumes and space guarantees, see the *Storage Management Guide*.

SnapDrive for Windows upgrade to version 3.1.1

Versions of SnapDrive™ for Windows earlier than 3.1.1 are incompatible with Data ONTAP 7.0.0.1 and higher. If you are upgrading to Data ONTAP 7.1 and you are using an earlier SnapDrive version, you must upgrade SnapDrive for Windows to version 3.1.1 or higher.

You are advised to upgrade SnapDrive on your Windows host before upgrading your storage system to Data ONTAP 7.1.

To download SnapDrive for Windows software from the NOW site to your Windows host, complete the following steps.

Step	Action
1	Use a Web browser to go to http://now.netapp.com/ .
2	Click Service & Support.
3	Click Download Software.

Step	Action
4	In the Select Platform list box in the SnapDrive product row, select Windows and click Go. Result: You see a list of the versions of SnapDrive for Windows. Click the 3.1.1 version link to see more information and to download software.

Change in CIFS SMB signing default behavior

By default, SMB signing is disabled in Data ONTAP 7.0.2 and later releases; it was enabled by default in Data ONTAP 7.0.1. If you are upgrading to Data ONTAP 7.1 from 7.0.1 and if you have configured CIFS clients in your environment to require SMB signing, you must enable SMB signing by setting the `cifs.signing.enable` option to on.

For more information, see the NetApp options(1) man page and the *File Access and Protocols Management Guide*.

Note

Enabling SMB signing in a Windows environment can affect the performance of both Windows clients and storage systems running Data ONTAP. For more information, see the *File Access and Protocols Management Guide*.

Data ONTAP 7.1 upgrades could take longer than usual

When you upgrade to Data ONTAP 7.1, Inode To Pathname (i2p) information is written to each existing volume. Depending on the number of volumes on your system, this could cause the upgrade to take more time than previous Data ONTAP upgrades to complete.

Recommendations for all storage systems upgrading to Data ONTAP 7.1

Keep the following tips in mind before upgrading to Data ONTAP 7.1:

- ◆ Upgrade during non-peak hours.
- ◆ Avoid performing a quota initialization prior to upgrading. If a quota initialization is begun prior to upgrading, wait for the initialization to complete.

Note

Beginning in Data ONTAP 7.0, quota initialization is terminated when you halt, reboot, or upgrade your storage system. For more information, see the *Storage Management Guide*.

Upgrading from releases earlier than Data ONTAP 7.0

Issues you must resolve before upgrading

This section discusses upgrade issues that you should be aware of if you are upgrading from a version of Data ONTAP earlier than 7.1:

- ◆ [“Important caution for storage systems with DS14 or DS14Mk2 disk shelves”](#) on page 24
- ◆ [“Upgrade recommendation for diskette-based storage systems before upgrading to Data ONTAP 7.0”](#) on page 26
- ◆ [“Upgrade recommendation for storage systems with FC-AL adapters”](#) on page 26

New behavior in upgraded systems

The following sections discuss new behavior in systems that have been upgraded to the Data ONTAP 7.0 release family:

- ◆ [“About volume types after upgrading to Data ONTAP 7.0 or later”](#) on page 28
- ◆ [“Space reservations changes for traditional volumes”](#) on page 28
- ◆ [“Changes to automatic editing for /etc/exports”](#) on page 29
- ◆ [“Transferring files larger than 4 TB”](#) on page 29
- ◆ [“Managing administrative access privileges after upgrading”](#) on page 30
- ◆ [“User name case-sensitivity”](#) on page 30
- ◆ [“Changes to vol status command output”](#) on page 31
- ◆ [“Changes to vol commands”](#) on page 31
- ◆ [“Quotas off during upgrade”](#) on page 32
- ◆ [“SSL must be re-enabled after upgrade”](#) on page 32
- ◆ [“Virus scan default timeout option”](#) on page 32

Important caution for storage systems with DS14 or DS14Mk2 disk shelves

If you are upgrading Data ONTAP software in systems with DS14 or DS14Mk2 disk shelves, including FAS200 series systems, and you are currently running a version of Data ONTAP that does not include a fix for Bug ID 113823 (Disk flash corrupted during firmware update), *do not* update disk firmware before upgrading Data ONTAP. Instead, upgrade your Data ONTAP software to a release that contains a fix for Bug ID 113823 first, then update disk firmware.

Attention

If your system is susceptible to Bug ID 113823, failure to upgrade software before updating disk firmware could cause disk drive damage and data loss.

For a list of Data ONTAP versions in which Bug ID 113823 has been fixed, see the following URL:

<http://now.netapp.com/NOW/cgi-bin/bugrellist?bugno=113823>

About Bug ID 113823 and its fix: Disk firmware is automatically updated on reboot as part of the Data ONTAP upgrade whenever the disk drives are running firmware older than the firmware version that is included with the new Data ONTAP release. As disk firmware is updated, it is possible (though unlikely) that all disk drives on a shelf will receive the update at the same time. This simultaneous update might cause the firmware update process to fail, rendering the disks unusable and causing loss of data if it happens to multiple disks in the same Redundant Array of Independent Disks (RAID) group.

The Data ONTAP 7.0 release family fixes Bug ID 113823 by preventing firmware from being concurrently updated on all disk drives. For more information about this issue, see the description for Bug ID 113823, on the NOW site at the following URL:

<http://now.netapp.com/NOW/cgi-bin/bol?Type=Detail&Display=113823>

Note

Bug ID 113823 occurs only when every disk drive on the shelf requires an update.

Checking disk firmware versions: If you have AutoSupport enabled, you can use the AutoSupport analysis tool, available on the NOW My Installed Products page, to verify your disk firmware version in advance and find out whether your system is vulnerable to Bug ID 113823.

Attention

If the AutoSupport analysis indicates that all disks on a disk shelf need firmware updates, failure to upgrade software before updating disk firmware could cause disk drive damage and data loss.

For instructions about using the AutoSupport analysis tool to detect disk firmware requirements, see “[Detecting outdated disk firmware](#)” on page 71.

**Upgrade
recommendation
for diskette-based
storage systems
before upgrading to
Data ONTAP 7.0**

If you are upgrading a diskette-based storage system from a version of Data ONTAP earlier than 6.2, you are advised to upgrade to Data ONTAP 6.3.3 or the latest release in the 6.3 release family first before upgrading to the Data ONTAP 7.0 release family.

By upgrading in two steps, you reduce the number of RAID and WAFL® (Write Anywhere File Layout) changes. As a result, the upgrade from Data ONTAP 6.3.3 to 7.0 should finish sooner with less risk of problems.

Systems affected: You are advised to upgrade the following storage systems to Data ONTAP 6.3.3 or the latest release in the 6.3 family before upgrading to the Data ONTAP 7.0 release family:

- ◆ F810
- ◆ F820
- ◆ F840

You can download Data ONTAP 6.3.3 from the Download Software page (on the NOW site at <http://now.netapp.com/NOW/cgi-bin/software/>). Under Data ONTAP, select your platform, then select View & Download under the Data ONTAP 6.3.3 release and follow the instructions.

**Upgrade
recommendation
for storage systems
with FC-AL
adapters**

When you upgrade to Data ONTAP 7.0 or later, you are advised to configure your Fibre Channel-Arbitrated Loop (FC-AL) adapters to run in 1-Gb and 2-Gb mode by changing the jumpers on the adapter. If you have not already done so, you are also advised to upgrade system firmware to Open Firmware (OFW) 4.3 or later in order to be compatible with the new 1-Gb and 2-Gb setting.

Note

Do not apply any configuration changes to FCP target-mode adapters.

FC-AL adapters configured for 1-Gb mode will continue to work in a Data ONTAP 7.1 environment without resetting the jumpers or upgrading firmware. However, 1-Gb and 2-Gb mode is recommended for best performance of these adapters.

Complete this procedure if your FC-AL adapters have the following marketing and spare part numbers:

- ◆ X2050A (SP-2050A)—Connects to storage
- ◆ X2051A (SP-2051A)—Connects to a tape backup device

To reconfigure your FC-AL adapters and upgrade firmware, complete the following steps.

Step	Action
1	Upgrade the system firmware to OFW 4.3 or later. You can find system firmware on the Service Image page (on the NOW site at http://now.netapp.com/NOW/download/tools/serviceimage/). Follow the instructions given there to download and install system firmware 4.3 or later. For more information on upgrading system firmware, see Chapter 3, “Upgrading NetApp Firmware,” on page 59. Note_____ If you are upgrading from 4.2.2 or 4.2.3 to 4.3 firmware, you must run the <code>set - defaults</code> command from the firmware ok prompt <i>before</i> rebooting your storage system. For more information, see Bug ID105365 (on the NOW site at http://now.netapp.com/NOW/cgi-bin/bol?Type=Detail&Display=105365).
2	Upgrade Data ONTAP to 7.0 or later as described in Chapter 4, “Upgrading Data ONTAP Software,” on page 73.
3	Remove any FC-AL adapters from your system chassis and verify that jumpers J3 and J4 are set in the 1-2 position as described in “Configuring new and existing FC-AL adapters” (on the NOW site at http://now.netapp.com/NOW/knowledge/docs/hardware/filer/adapter.pdf). Reposition the jumpers if necessary. Note_____ Consult your hardware storage system guide before opening the system and removing adapters.

Systems affected: FC-AL adapters can be installed on the following storage systems:

- ◆ R200
- ◆ R150
- ◆ FAS960
- ◆ FAS940
- ◆ F880
- ◆ F825

About volume types after upgrading to Data ONTAP 7.0 or later

When you upgrade to Data ONTAP 7.0 or later, any existing volumes are preserved as traditional volumes. Both traditional and FlexVol volumes are supported in Data ONTAP 7.1; you are not required to use FlexVol volumes.

If you want to take advantage of FlexVol volume technology, you must create new FlexVol volumes and aggregates, then migrate your data from existing traditional volumes to the new FlexVol volumes using the `ndmcopy` command.

Note _____
Snapshot copies on the original volume are unaffected by the migration, but they are not valid for the new volume.

For more information about...	See the...
Migrating data to FlexVol volumes	<i>Storage Management Guide</i>
Using the <code>ndmcopy</code> command to migrate data	<i>Data Protection Tape Backup and Recovery Guide</i>
Planning and configuring aggregates and volumes	<i>Storage Management Guide</i>

Note _____
In the Data ONTAP 7.0 release family, the storage system root volume can be either a traditional or a FlexVol volume. However, because you have more control over the size of FlexVol volumes than you do over traditional volumes, changing the traditional root volume to a flexible root volume allows you to organize your storage more efficiently. For information about changing the root volume to a FlexVol volume, see the *System Administration Guide*.

Space reservations changes for traditional volumes

To provide space reservations for traditional volumes in the Data ONTAP 7.0 release family, you must assign file reservation after each file is created.

If your system includes volumes with default file space reservation enabled (using the `create_reserved` parameter of the `vol options` command), default space reservation is no longer available in the corresponding traditional volume after upgrading to the Data ONTAP 7.0 release family. Any files created in the volume after upgrading will not have space reservations enabled unless specified on a per-file basis. However, space reservations will remain enabled for any files created *before* upgrading.

For more information about space reservation and guarantees, see the *Storage Management Guide*.

Changes to automatic editing for /etc/exports

Beginning with the Data ONTAP 6.5 release family, several `vol` command options cause automatic updates to the `/etc/exports` file, as described in “[Automatic editing of the /etc/exports file](#)” on page 40.

In the Data ONTAP 7.0 release family, the following additional operations cause the exports file to be edited:

- ◆ `vol clone`
When a clone is created from a FlexVol volume, a new clone volume is created in the exports file. The cloned volume has the same export rules as the parent.
- ◆ `vfiler migrate`
If a vFiler™ unit is migrated to a new physical storage system, the exports file is modified to reflect any changes to volume configuration.

Transferring files larger than 4 TB

The Data ONTAP 7.0 release family supports files with up to 16 TB of data. However, transfers of data to systems running earlier versions of Data ONTAP fail if the file system being transferred contains a file larger than 4 TB. The following features can be affected by this issue.

- ◆ SnapMirror for qtree replication
A qtree replication from a SnapMirror source system running Data ONTAP 7.0 or higher will fail if any mirrored qtree includes a file larger than 4 TB and the qtree SnapMirror destination system is running a version of Data ONTAP earlier than 7.0.
- ◆ SnapVault
SnapVault transfers from a primary storage system fail if the qtree or directory to be backed up includes a file larger than 4 TB and the SnapVault secondary system is running a version of Data ONTAP earlier than 7.0.
- ◆ The restore command
Data restorations from a tape created with the `dump` command on a storage system running Data ONTAP 7.0 or later fail if the tape archive includes a file larger than 4 TB and the restoration target system is running a version of Data ONTAP earlier than 7.0.

Note

It is not a requirement for qtree SnapMirror destination systems, SnapVault secondary systems, or restoration target systems to run a version of Data ONTAP that is the same or later than the version the source system is running. The requirement to upgrade destination systems before source system applies only to SnapMirror for volume replication. Nonetheless, if you upgrade systems with these configurations to Data ONTAP 7.0 or later before upgrading the corresponding source systems, you will not encounter problems transferring files up to 16 TB.

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

For more information about the restore command, see the *Data Protection Tape Backup and Recovery Guide*.

Managing administrative access privileges after upgrading

The Data ONTAP 7.0 release family provides enhanced capabilities for access control and administrative account management. When you upgrade to Data ONTAP 7.0 or later, any existing local users who are not root users are granted full privileges. To manage or reduce those privileges, you must create one or more groups with defined roles and capabilities, then place the users in groups.

For more information about user account management, see the *System Administration Guide*.

User name case-sensitivity

In the Data ONTAP 7.0 release family, user names are no longer case-sensitive. For example, in previous releases the names “admin” and “Admin” can stand for two different users, but in Data ONTAP 7.0 and later, the second name is a duplicate of the first.

When you upgrade to Data ONTAP 7.0 or later, account names are scanned, an error message is displayed if such duplicate user names are detected, and the second name is given a suffix; for example, the duplicate “Admin” would become “admin-a”.

In this case, any user attempting to log in to the storage system with the Admin account name and password would be denied. To correct this, you must give the user the modified account name or assign a new account name with appropriate privileges.

For more information about user account management, see the *System Administration Guide*.

Changes to vol status command output

The `vol status` command has been modified in the Data ONTAP 7.0 release family to distinguish between flexible and traditional volumes. If you created scripts in earlier releases, some command output has changed, which might affect the behavior of your script.

Specifically, the Status column of the `vol status` command output is now a comma-separated list including RAID type and storage container type (flexible or traditional), as shown in the following example.

```
filer12> vol status
      Volume State      Status      Options
      tv online        raid4, trad resyncsnaptime=10
      mtv online        raid4, trad
                        mirrored
      vol0 online       raid_dp, flex  root
      fv1 online       raid_dp, flex
      fv2 offline       raid_dp, flex
```

For more information about the `vol status` command and output, see the `vol(1)` manual page.

Changes to vol commands

Except as noted, the `vol` commands continue to function as they have in previous releases. However, the Data ONTAP 7.0 release family introduces the `aggr` command set, which is used to manage the physical layer of storage management presented in aggregates. If you use a `vol` command to manage physical storage components, such as disks or RAID, you will see an advisory error message recommending that you use the appropriate `aggr` command.

For example, if you use the `vol mirror` command to mirror a FlexVol volume, you will see the following screen message:

```
filer12> vol mirror fv1 -d v0.25 -dv0.26
vol mirror: 'fv1' is a flexible volume.
Consider using the 'vol container fv1' command to find its
containing aggregate, then 'aggr mirror' to mirror it.
```

For more information about aggregate and volume management commands, see the follow documents:

- ◆ `aggr(1)` man page
- ◆ `vol(1)` man page
- ◆ *Storage Management Guide*

Note

You are advised to use the `aggr` command to manage physical storage components.

Quotas off during upgrade

Quotas are turned off during the upgrade to the Data ONTAP 7.0 release family. This is because the online quota format is changed in this release and the quota database must be rebuilt. As part of the upgrade, quotas are automatically turned back on when the database has been rebuilt and the upgrade is complete.

If you have set quotas, and you want to verify that the quota database has been rebuilt and quotas are back on, enter the `quota` command without options. For more information about using the command, see the `quota(1)` man page.

SSL must be re-enabled after upgrade

The Data ONTAP 7.0 release family includes built-in SecureAdmin™ functionality with SSL (Secure Socket Layer) protocol. It is no longer necessary to install and license it as a separate product.

If you installed and licensed the SecureAdmin product in earlier Data ONTAP releases, your existing SSL certificates are converted on upgrade. However, you must re-enable SSL after the upgrade for the new certificates to take effect. To do so, enter the following command after the upgrade has completed:

```
secureadmin enable ssl
```

For more information about SSL and SecureAdmin functionality in the Data ONTAP 7.0 release family, see the *System Administration Guide*.

Virus scan default timeout option

Beginning with Data ONTAP 6.4.4, the `vscan timeout` option had a default setting of 10 seconds. If your storage system has been upgraded at any time from Data ONTAP 6.4.3 or earlier and you have not explicitly reset this option, the original default value of 30 seconds is still in effect.

This option controls the length of the waiting period before the scanning client checks the status of a virus scan request. By setting the option to a lower value, you enable the storage system to determine more quickly if there is a problem with a virus scan request. You are therefore advised to set this option to 10 seconds or another lower value appropriate to your configuration.

To view the current setting of the `vscan timeout` option, enter the following command at the storage system console:

```
vscan options timeout
```

To set the `vscan timeout` option to a new value, enter the following command at the storage system console (where *nn* is the new value in seconds):

```
vscan options timeout set nn
```

For more information about virus scanning in the Data ONTAP 7.0 release family, see the *Data Protection Online Backup and Recovery Guide* and the `vscan(1)` man page.

Upgrading from releases earlier than Data ONTAP 6.5

Information about the Data ONTAP 6.5 release family you should know

This section discusses upgrade issues that you should be aware of if you are upgrading from releases earlier than Data ONTAP 6.5.

Issues you must resolve before upgrading

The following sections discuss upgrade issues that you must resolve before upgrading to the Data ONTAP 7.1 release family from releases earlier than Data ONTAP 6.5:

- ◆ [“Storage systems with CompactFlash cards that use OFW must be running firmware 4.2.2 or later”](#) on page 35
- ◆ [“Important caution for storage systems with CompactFlash cards running Data ONTAP 6.4.1”](#) on page 36
- ◆ [“Storage Systems with CompactFlash cards must be running the latest version of Data ONTAP 6.4.x before upgrading to Data ONTAP 6.5 or later”](#) on page 36
- ◆ [“Open Systems SnapVault Agent upgrade to version 1.1”](#) on page 37

New behavior in upgraded systems

The following sections discuss new behavior in systems that have been upgraded to the Data ONTAP 7.1 release family from releases earlier than Data ONTAP 6.5:

- ◆ [“Increased disk and CPU activity in Data ONTAP 6.5 and later”](#) on page 39
- ◆ [“New access rules for /etc/exports file”](#) on page 39
- ◆ [“Automatic editing of the /etc/exports file”](#) on page 40
- ◆ [“RAID4 groups size limits”](#) on page 40
- ◆ [“Upgrading oversized RAID4 groups to RAID-DP”](#) on page 41
- ◆ [“iSCSI CHAP configuration not saved on upgrade”](#) on page 43
- ◆ [“Resetting the NFS exports timeout option value”](#) on page 43
- ◆ [“Default behavior change for FTP”](#) on page 44
- ◆ [“Creating vFiler shares using the MMC Computer Management tool”](#) on page 44

Storage systems with CompactFlash cards that use OFW must be running firmware 4.2.2 or later

For storage systems that have CompactFlash cards and that run Open Firmware (OFW), you must upgrade the system firmware and diagnostics files to version 4.2.2 or later before upgrading to the Data ONTAP 7.0 release family. Data ONTAP images for release later than 6.5 will not fit on the 32-MB CompactFlash cards as they are currently formatted. If you attempt to download Data ONTAP 6.5 or later on your storage system before upgrading to firmware version 4.2.2 or later, the download will stop and the storage system will display a message similar to the following:

```
download: Failed to repartition the card.
download: The current firmware version does not support
repartitioning of the boot device as required.
download: Please update your system firmware to 4.2.2 or later
prior to upgrading to this OS. The System Firmware area on the NOW
site has further information.
```

Note

If you are running an OFW release earlier than 4.2.2, make sure that you upgrade to firmware version 4.3 or later. Do not install firmware version 4.2.2 or 4.2.3.

If you are running OFW release 4.2.2 or 4.2.3, make sure that you upgrade to firmware version 4.3 or later to correct Bug ID 105365. This bug can prevent the system from booting after running the `update-flash` command. For more information, see Bug ID 105365 (on the NOW site at the following URL: <http://now.netapp.com/NOW/cgi-bin/bol?Type=Detail&Display=105365>).

Considerations concerning system firmware upgrades: Review the following issues before upgrading system firmware.

- ◆ If you are running Data ONTAP 6.4.1, see “[Important caution for storage systems with CompactFlash cards running Data ONTAP 6.4.1](#)” on page 36, *before* you upgrade your system firmware and diagnostics to version 4.3 or later.
- ◆ If you are upgrading from 4.2.2 or 4.2.3 to 4.3 or later system firmware, you must run the `set-defaults` command from the firmware ok prompt *before* rebooting. For more information, see Bug ID 105365.
- ◆ If you are upgrading to system firmware version 4.3 or later from a version other than 4.2.2 or 4.2.3, or if you are already running system firmware version 4.3 or later, do not run the `set-defaults` command.

Systems affected: The following storage systems must be running firmware version 4.2.2 or later before upgrading to the Data ONTAP 7.0 release family:

- ◆ F825
- ◆ F880

- ◆ FAS940
- ◆ FAS960
- ◆ R100
- ◆ R150

You can find system firmware on the NOW site at <http://now.netapp.com/NOW/download/tools/serviceimage/>. Follow the instructions there to download and install the latest system firmware.

Important caution for storage systems with CompactFlash cards running Data ONTAP 6.4.1

If you are running Data ONTAP 6.4.1 on a storage system that has a CompactFlash card and that runs OFW, you must complete an additional procedure to manually format the service partition before upgrading the system firmware to version 4.3 or later. If you do not complete this procedure, your storage system will be unable to upgrade to 4.3 and future firmware releases. For more information, see Bug ID 98952 on the NOW site at the following URL: <http://now.netapp.com/NOW/cgi-bin/bol?Type=Detail&Display=98952>

Systems affected: This caution applies to the following storage systems running Data ONTAP 6.4.1 before upgrading to the Data ONTAP 7.0 release family:

- ◆ F825
- ◆ F880
- ◆ FAS940
- ◆ FAS960
- ◆ R100
- ◆ R150

You can find the NetApp OFW/Diagnostics Download Instructions for CompactFlash Units and firmware/diagnostics images on the NOW site at <http://now.netapp.com/NOW/download/tools/serviceimage/30800128/>.

Storage Systems with CompactFlash cards must be running the latest version of Data ONTAP 6.4.x before upgrading to Data ONTAP 6.5 or later

If you are running a version of Data ONTAP earlier than 6.4.2 on storage systems that have a CompactFlash card, you must install Data ONTAP 6.4.2 or the latest Data ONTAP release in the 6.4 family, before upgrading to the Data ONTAP 7.0 release family. If you attempt to download Data ONTAP 6.5 or later on your storage system before installing the latest version in the Data ONTAP 6.4 family, the download will stop and the storage system will display a message similar to the following:

download: The boot image is larger than the primary boot device (23068672 bytes instead of 11534336).
download: Reformatting boot device.
download: Unable to reformat the boot device. download failed.
Cannot boot off disks with latest image

Systems affected: The following storage systems must be running Data ONTAP 6.4.2 or later before upgrading to the Data ONTAP 7.0 release family:

- ◆ F87
- ◆ F825
- ◆ F880
- ◆ FAS940
- ◆ FAS960
- ◆ R100
- ◆ R150

You can download the latest release in the Data ONTAP 6.4 family from the Download Software page (on the NOW site at <http://now.netapp.com/NOW/cgi-bin/software/>). Under Data ONTAP, select your platform, then select View & Download under the latest Data ONTAP 6.4 family release and follow the instructions.

To upgrade your storage system software, follow the instructions specific to your storage system in Chapter 4, “[Upgrading Data ONTAP Software](#),” on page 73.

Open Systems SnapVault Agent upgrade to version 1.1

Open Systems SnapVault Agent 1.0 is incompatible with Data ONTAP 6.5 and higher. If you are upgrading to the Data ONTAP 7.0 release family and you are using Open Systems SnapVault Agent 1.0, you must upgrade Open Systems SnapVault Agent to version 1.1 or higher.

You are advised to upgrade the Open Systems SnapVault agent on your host platform *before* upgrading your storage system to the Data ONTAP 7.0 release family.

To download Open Systems SnapVault agent software from the NOW site to your host platform, complete the following steps.

Step	Action
1	Use a Web browser to go to http://now.netapp.com/ .
2	Click Service & Support.

Step	Action
3	Click Download Software.
4	In the Select Platform list box in the Open Systems SnapVault Agent product row, select your host platform type.
5	Click Go. Result: You see a list of the versions of Open Systems SnapVault agent software supported by your platform host. Click on the version links for more information and to download software.

**Recommendations
for all storage
systems upgrading
from Data ONTAP
releases earlier
than 6.5**

The Data ONTAP 6.5 release family introduced support for larger inodes. For the smoothest conversion to larger inodes, keep the following tips in mind before upgrading to the Data ONTAP 7.0 release family from a release earlier than 6.5:

- ◆ To complete the upgrade more quickly, upgrade during non-peak hours and keep storage system workload moderate.
- ◆ If you are serving files with the CIFS protocol, make sure the file system is not 100 percent full.
- ◆ Avoid performing a quota initialization prior to upgrading. If a quota initialization is begun prior to upgrading, wait for the initialization to complete.

Because of large inode support, upgrading to Data ONTAP 6.5 or later can take significantly longer to complete than upgrades to releases earlier than 6.5. Depending on a number of factors—particularly hardware platform, total number of inodes, and percentage of inodes in use—such an upgrade can take anywhere from a few minutes to several hours. For example, upgrading a FAS980 system with 95.6 million inodes, 73% of which are in use, takes around 67 minutes to complete. (These figures represent guidelines based on NetApp customer configurations. Your system upgrade time could differ based on system load and other factors in your environment.)

Note

During the upgrade, the storage system will continue to serve data. Data ONTAP upgrades takes place in the background, so the system is completely available during the upgrade.

Increased disk and CPU activity in Data ONTAP 6.5 and later

After upgrading to the Data ONTAP 7.0 release family from a release earlier than 6.5, you might observe disk LEDs blinking on an apparently idle system. You might also observe some CPU activity even when there is no user workload present. This is due to a default “media scrub” background process that continuously finds and cleans media errors on the storage system disks. This process is designed to work in self-throttling mode with negligible impact on storage system performance.

New access rules for `/etc/exports` file

Beginning with the Data ONTAP 6.5 release family, access rules in the `/etc/exports` file were changed. Rules that previously allowed read access to all hosts have been changed to allow no access. For example, this is a rule in versions of Data ONTAP earlier than 6.5:

```
/vol/vol1 -rw=host1:host2
```

The rule meant that the hosts named `host1` and `host2` both had read-write access to `vol1`, while all other hosts had *read access*.

In Data ONTAP 6.5 and later releases, however, the rule now means that `host1` and `host2` have read-write access and all other hosts have *no access*. In order to allow read access for all hosts in Data ONTAP 6.5 and later, you need to add `-ro` after `vol1` to allow read access to all hosts:

```
/vol/vol1 -ro,rw=host1:host2
```

The first time you upgrade to the Data ONTAP 7.0 release family from a release earlier than 6.5, the `/etc/exports` file is changed automatically to be sure your rules work the same way. The rule `/vol/vol1 -rw=host1:host2` is converted to the form `/vol/vol1 -ro,rw=host1:host2`.

When subsequent upgrades are performed, Data ONTAP assumes that the rules in the `/etc/exports` file have already been converted and additional conversions do not take place.

Other changes to the `/etc/exports` file in Data ONTAP 6.5 and later releases include the following:

- ◆ Line length is no longer restricted.
- ◆ The `-access` option is no longer available.

The `-access` option was used to list hosts that could have any access. Instead, the `-rw` and `-ro` options are used together to enable the same privileges.

- ◆ You can now list netgroups in the `rw=`, `ro=`, and `root=` options.
- ◆ The number of entries in an access option is no longer restricted.

For more information about changes to the `/etc/export` file upon upgrade, see the `exports(5)` man page.

Automatic editing of the `/etc/exports` file

In Data ONTAP 6.5 and later releases, the following operations cause the `/etc/exports` file to be modified according to the new rules described in “[New access rules for `/etc/exports` file](#)” on page 39:

- ◆ `vol create`

New volumes are automatically entered in `/etc/exports`.

Note

When new volumes are created using the `vol create` command, access rules for the new volumes are set by default to `-rw,nosuid` (read-write permissions, `setuid` and `setgid` not allowed) in `/etc/exports`. If you want different access permissions, you must edit `/etc/exports` after the `vol create` operation is complete.

- ◆ `vol destroy`

Volume listings are removed from `/etc/exports` when the volume is destroyed.

- ◆ `vol rename`

When a volume is renamed, the new volume name is automatically entered in `/etc/exports`.

RAID4 groups size limits

The maximum size limits for RAID4 groups changed in the Data ONTAP 6.5 release family. The `raidsize` option of the `vol` command controls the size of new RAID groups that might be created in the volume. When new volumes are configured with RAID4 protection, they are created according to the `raidsize` settings for your platform.

If you upgrade to Data ONTAP 6.5 or later from a release earlier than 6.5 on a system with RAID4 protection, and if the settings for existing volumes are higher than the new limits, the `raidsize` option for those volumes is changed to the new RAID4 group size limit for your model appliance. If the existing volume’s `raidsize` option is already set below the new limits, it is not changed.

These are the default and maximum RAID4 group sizes for NetApp® appliances affected by this change:

- ◆ The R100 default and limit is 8.
- ◆ The R150 default and limit is 6.
- ◆ The FAS250 default is 7, and the limit is 14.
- ◆ The default for other storage systems is 8, and the limit is 14.

The sizes of existing volumes' RAID4 groups remain the same as before the upgrade, even if the size exceeds the new limits. However, if the `raidsz` option changes to the new, reduced limit, any new RAID4 groups that are created as disks are added to the volumes will conform to the new `raidsz` option.

Example: An R150 running Data ONTAP 6.4 has a volume RAID group size of 8. When the appliance is upgraded to Data ONTAP 6.5 or later, `raidsz` is changed to 6, the new limit. Existing RAID groups can have up to 8 disks, but any new RAID groups created in that volume will have up to 6 disks, which is also the default.

Note

Default and maximum `raidsz` values might change in future Data ONTAP releases. For a listing of current values, see the *Storage Management Guide*.

Upgrading oversized RAID4 groups to RAID-DP

After the upgrade to the Data ONTAP 7.0 release family from a release earlier than 6.5, if you have volumes with RAID4 groups whose size exceeds the new RAID4 limit, you are strongly advised to upgrade to RAID-DP protection for those volumes. For volumes with oversized RAID groups, Data ONTAP issues a daily message similar to the following:

```
Wed Oct 1 03:44:51 GMT [raid.config.rg.size.reminder:warning]:
vol1 has a raid group larger than the limit for raid4 raidtype
(max. raid group size: 7, limit is 6). Convert the volume to
raid_dp.
```

Note

Converting from RAID4 to RAID-DP could take one or more hours depending on the volume size and underlying RAID groups. During this time, performance might be degraded. You are therefore advised to schedule conversions from RAID4 to RAID-DP during non-peak hours.

To convert the RAID protection from RAID4 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 volume by entering the following command. <pre>vol status volname -r</pre> <i>volname</i> is the name of the oversized RAID4 volume.
2	Enter the following command to make sure that a hot spare disk exists on the storage system for each RAID group listed for the volume, and make sure that the hot spare disks are equal in size to the existing parity disks in those RAID groups. <pre>vol status -s</pre> If necessary, add the appropriate number of disks of the appropriate size and checksum type to the spare pool.
3	Convert the RAID protection for the volume in question by entering the following command: <pre>vol options volname raidtype raid_dp</pre> <i>volname</i> is the name of the volume you want to convert. Result: Data ONTAP converts the specified volume's RAID protection from RAID4 to RAID-DP. It adds a disk from the hot spare disk pool to each existing RAID group in the volume and makes that disk the additional "dparity" disk to provide double parity protection. Any subsequent RAID groups that might be created as disks are added to the specified volume will conform to the new RAID-DP group size defaults and limits for the storage system in question.

Default and maximum RAID-DP group sizes for the Data ONTAP 7.0 release family are as follows:

- ◆ The R100 default and limit is 12.
- ◆ The R150 default and limit is 12.
- ◆ The R200 default and limit is 14.
- ◆ The default for other storage systems is 16, and the limit is 28.

For more information, see the *Storage Management Guide*.

iSCSI CHAP configuration not saved on upgrade

If you upgrade from the Data ONTAP 6.4 release family and you have Challenge Handshake Authentication Protocol (CHAP) configured for Internet SCSI (iSCSI), the CHAP configuration from the previous release is not saved.

The CHAP configuration file in Data ONTAP 6.5 and later uses a new format that is not compatible with the CHAP configuration file format of previous releases. When you upgrade, you must use the `iscsi security` command to reconfigure CHAP settings.

For information about using the `iscsi security` command, see the *Block Access Management Guide*.

Resetting the NFS exports timeout option value

Data ONTAP 6.5.2 and later releases include a fix for bug 129083, which made it possible for NFS exports to be improperly rendered inaccessible after Data ONTAP is restarted.

If you modified the value of the `nfs.export.neg.timeout` option as a workaround for this problem in an earlier Data ONTAP 6.5 family release, you can reset this option to its default value (3600), or any other value you used for this option before applying the workaround, after you upgrade to the Data ONTAP 7.0 release family.

Leaving the lowered (workaround) option value may cause an increase in name-service traffic when several clients attempt to access resources for which they do not have permissions.

To restore the default option value for caching of denied NFS requests, complete the following step:

Step	Action
1	From the Data ONTAP command line, enter the following command: <pre>options nfs.export.neg.timeout 3600</pre> Result: This command sets the length of time (in seconds) for which denied access requests are cached. Once the cache is cleared and the name service is active, access requests will be handled properly.

Default behavior change for FTP

An FTP security enhancement option (`ftp.3way.enable`) was added that changes the default behavior of Data ONTAP beginning with the 6.5 release family. This option prevents clients from redirecting FTP operations to a third-party client or server.

Standard FTP commands allow FTP communications to be redirected to third-party FTP servers. However, this functionality can be exploited using “bounce attacks” to circumvent network security and redirect FTP output (or FTP commands stored in a file) to unauthenticated machines.

In earlier versions of Data ONTAP: In earlier versions, standard FTP redirect features were enabled. This functionality raises security issues and has been removed from the default behavior of Data ONTAP.

In Data ONTAP 6.5 and later releases: In this version, FTP redirects have been disabled by default. You need to modify the value of the `ftp.3way.enable` option to permit FTP redirects.

To enable one client to use FTP to send a file directly from a storage system to another client, complete the following step.

Step	Action
1	From the Data ONTAP command line, enter following command: options ftp.3way.enable on Result: Redirecting of FTP files to third-party clients is immediately enabled.

Note

This option is persistent across storage system reboots. Valid values for `ftp.3way.enable` are `on` (allowing FTP redirects) and `off` (not allowing FTP redirects).

Creating vFiler shares using the MMC Computer Management tool

Beginning with Data ONTAP 6.5, there is a per-vFiler unit pseudo-root directory structure that allows all of a vFiler unit’s storage to be shared under a single share. However, that change eliminated special handling for `/vol/volname` for the `C$` share, which previously enabled the Microsoft Management Console (MMC) Computer Management tool to create the shares.

If you want to continue using the MMC Computer Management tool to create vFiler shares after upgrading to Data ONTAP 6.5 and later releases, you must delete any existing C\$ vFiler share and re-create it to point to the vFiler pseudo-root. For more information, see Bug ID 112857 (on the NOW site at the following URL:
<http://now.netapp.com/NOW/cgi-bin/bol?Type=Detail&Display=112857>).

Upgrading from releases earlier than Data ONTAP 6.4

Information about the Data ONTAP 6.4 release family you should know

The following sections discuss upgrade issues that you should be aware of if you are upgrading from releases earlier than Data ONTAP 6.4.

Macintosh resource.frk files affected

If your storage system is accessed by Macintosh computers using the DAVE™ client from Thursby Systems, the DAVE client starts using the new streams feature when Data ONTAP 6.4 or later is installed. Information that was stored in special resource.frk files is lost. To solve this problem, you should run the na_macstream conversion utility from a Windows client after the upgrade is completed and before DAVE clients access any files. For instructions about obtaining, installing, and running the na_macstream utility, see Appendix B, “[DAVE Macintosh Client resource.frk Conversion](#),” on page 167.

Space reservation settings

If you are using space reservations with Data ONTAP 6.3, space reservations on volumes are ignored when you upgrade to later versions of Data ONTAP. Space reservations are turned on for all logical unit numbers (LUNs) and virtual local disks (VLDs), but are left off for all other files.

Upgrading from releases earlier than Data ONTAP 6.2

Information about the Data ONTAP 6.2 release family you should know

The following sections discuss upgrade issues that you should be aware of if you are upgrading from releases earlier than Data ONTAP 6.2.

New options and settings for `raid.reconstruct_speed`

If you are upgrading from a release that has the `raid.reconstruct_speed` feature, you should be aware that the options and settings for this feature changed in Data ONTAP 6.2. The `raid.reconstruct_speed` option can still be set, but it only affects the value of `raid.reconstruct.perf_impact`, which is one of four new options. The ten settings previously used for the `raid.reconstruct_speed` feature translate to three settings for the new options, as follows:

- ◆ 7-10 = high
- ◆ 5-6 = medium
- ◆ 1-4 = low

For more information about changing the RAID data reconstruction speed, see the latest *Storage Management Guide*.

Change in vif behavior

Load-balancing algorithms for virtual interfaces (vifs) were introduced in Data ONTAP 6.2. If you are using vifs, and you upgrade to Data ONTAP 6.2 or later from an earlier release, the network traffic flow out of the system might change. This could be because the vifs follow the default IP-based load-balancing policy. However, since the default load-balancing policy for most of the switches is IP-based, traffic flow is usually not disturbed. If the outbound flow of the network is not sufficiently balanced, you can change the load-balancing policy on the vif. For more information, see the *Network Management Guide*.

Upgrading from releases earlier than Data ONTAP 6.1

Information about the Data ONTAP 6.1 release family you should know

The following sections discuss upgrade issues that you should be aware of if you are upgrading from releases earlier than Data ONTAP 6.1.

Applying the same quota to a user's UNIX and Windows IDs

In Data ONTAP 6.1 and later, both UNIX and Windows IDs can be specified in the `/etc/quotas` file to track user quotas. If you previously specified a quota based on a user's UNIX IDs, and then after upgrading you specify a quota for the same user based on the Windows ID, these are treated as separate quotas for the user. To apply the same quota whether the user logs in with the UNIX ID or Windows ID, add the following line to the `/etc/quotas` file:

QUOTA_PERFORM_USER_MAPPING on

For more information about disk space management using quotas, see the *Storage Management Guide*.

Restarting quotas

Because of the addition of Security ID (SID)-based quotas in Data ONTAP 6.1, you might see a warning message about correcting quotas if you upgrade from an earlier version of Data ONTAP.

If you see a warning message, restart quotas by completing the following steps.

Step	Action
1	Turn quotas off using the following command: quota off volume
2	Turn quotas back on using the following command: quota on volume

**Earlier snapshots
might need to be
deleted**

If you are upgrading from a version of Data ONTAP earlier than 6.1, you should delete any snapshots taken on earlier releases due to changes made to the SnapMirror feature in Data ONTAP 6.1.

**Clustered storage
system disk
firmware upgrade
procedure changed**

The procedure for updating a cluster's disk firmware changed with Data ONTAP 6.1.1. For information about current disk firmware upgrade procedures, see “[Upgrading disk firmware](#)” on page 67.

**Local administrator
account
nonexistent**

After you upgrade from a version of Data ONTAP earlier than 6.1.1, when you reboot with CIFS enabled, you might see a message that the builtin\administrator account does not exist. Add a local user named “administrator” to resolve this issue. For information about adding a local user, see the *File Access Management Guide*.

**ApplianceWatch
upgrade**

Data ONTAP 6.1.1 and later is compatible with Appliance Watch™ 1.1 software for HP® OpenView®, and ApplianceWatch 1.0 for Tivoli® Enterprise™.

If you need to upgrade ApplianceWatch, you can find it on the Download Software page on the NOW site at the following URL:

<http://now.netapp.com/NOW/cgi-bin/software>

Upgrading iSCSI hosts

About target portal group tags in Data ONTAP 7.1

In Data ONTAP 7.1, the iSCSI target portal group tags changed from the values used in previous releases. These tags are used by iSCSI hosts to identify the iSCSI target provided by the storage system.

The Linux and HP-UX iSCSI software initiators create a persistent association between the target portal group tag value and the target. If the target portal group tag changes, the LUNs from the target become unavailable because the initiators recognize the LUNs as new devices.

Affected systems

All NetApp storage systems that provide iSCSI LUNs to hosts running any version of a Linux software initiator are affected. The following procedures are required when upgrading from any previous version of Data ONTAP to 7.1 or later.

Procedure overview

To upgrade Data ONTAP with iSCSI hosts running Linux or HP-UX, record the target portal group tags used by the storage system before the upgrade and then configure the interfaces to use the old tag values after the upgrade.

Attention

You must complete the first steps of this procedure *before* installing the new version of Data ONTAP software.

You must complete the following procedures in the order shown:

- ◆ Record existing portal group tags.
- ◆ Install the new Data ONTAP version but do not reboot the storage system.
- ◆ Stop the iSCSI service on the storage system.
- ◆ Reboot the storage system.
- ◆ Change the new portal group tags to the previous values.
- ◆ Start the iSCSI service on the storage system.

Recording existing portal group tags

Before upgrading, collect and record the portal group tags for each Ethernet interface on the storage system.

Collecting tags using the command line: To collect the existing portal group tag values using the command line, complete the following steps.

Step	Action
1	On the storage system console, enter the following command: iscsi show adapter Result: The current target portal groups are displayed. Example: <pre>filer2> iscsi show adapter Adapter: iswta Slot: N/A Description: NetApp Software Implementation Status: Online Target Portal Groups: portal group 1: inet 10.60.128.101 port 3260 portal group 2: inet 10.60.128.102 port 3260 portal group 3: inet 10.60.131.102 port 3260</pre>
2	Note the value associated with each IP address.

Step	Action
3	Enter the following command: <pre>ifconfig -a</pre> Result: Networking parameters, including the IP address, are displayed for each interface. Example: <pre>filer2> ifconfig -a e0: flags=848043<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500 inet 10.60.128.101 netmask 0xffffffff broadcast 10.60.128.255 ether 00:a0:98:01:25:0f (100tx-fd-up) e9a: flags=148043<UP,BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 1500 inet 10.60.128.102 netmask 0xffffffff broadcast 10.60.128.255 ether 00:07:e9:0d:8a:ce (1000fx-up) flowcontrol full e9b: flags=948043<UP,BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 1500 inet 10.60.131.102 netmask 0xffffffff broadcast 10.60.131.255 ether 00:07:e9:0d:8a:cf (1000fx-up) flowcontrol full</pre>
4	Note the interface identifier, such as e9a, associated with the IP addresses you collected in step 1, and then map to the portal group tags. Example: <pre>Interface e0 = portal group 1 Interface e9a = portal group 2 Interface e9b = portal group 3</pre>

Collecting tags using FilerView: To collect the existing portal group tag values using FilerView, complete the following steps:

Step	Action
1	Click LUNs > iSCSI > Report. Result: The reports page is displayed, including the iSCSI Adapters section at the bottom of the page. Example: <pre> Adapter: iswta Slot: N/A Description: NetApp Software Implementation Status: Online Target Portal Groups: portal group 1: inet 10.60.128.101 port 3260 portal group 2: inet 10.60.128.102 port 3260 portal group 3: inet 10.60.131.102 port 3260 </pre>
2	Note the value associated with each IP address.
3	Click Network > Manage Interfaces Result: The interfaces and their IP addresses are displayed.
4	Note the interface identifier, such as e9a, associated with the IP addresses you collected in step 1, and then map to the portal group tags. Example: <pre> Interface e0 = portal group 1 Interface e9a = portal group 2 Interface e9b = portal group 3 </pre>

Installing Data ONTAP

Follow the procedures in this guide to install the new version of Data ONTAP. Do not reboot the storage system yet.

Stopping the iSCSI service

Stop the iSCSI service on the storage system. This prevents the iSCSI hosts from seeing the new target portal group tags that are assigned by default when the Data ONTAP upgrade is complete.

Note

All iSCSI LUNs will be unavailable when you stop the iSCSI service.

Stopping the iSCSI service using the command line: To stop the iSCSI service on the storage system using the command line, complete the following step:

Step	Action
1	On the storage system console, enter the following command: iscsi stop

Stopping the iSCSI service using FilerView: To stop the iSCSI service on the storage system using FilerView, complete the following steps:

Step	Action
1	Click LUNs > Enable/Disable.
2	Clear the Enable check box for iSCSI.
3	Click Apply.

Rebooting the storage system

Once you have stopped the iSCSI service, you are now ready to reboot the storage system. The storage system reboots with the new version of Data ONTAP, but the iSCSI service is not enabled.

Changing the target portal group tags

The upgrade to the new version of Data ONTAP changed the target portal group tag values. You need to change the new values back to the values you recorded before the upgrade.

Data ONTAP 7.1 supports user-created target portal groups with user-defined tag values. You need to create a new target portal group for each interface and assign a new tag value to the group.

Changing target portal group tags using the command line: To change the target portal groups to their original values, complete the following steps:

Step	Action
1	View the new target portal groups and their tag values by entering the following command on the storage system console: iscsi tpgroup show Example: <pre>filer2> iscsi tpgroup show TPGTag Name Member Interfaces 1000 e0_default e0 1001 e9a_default e9a 1002 e9b_default e9b</pre>
2	Create a new target portal group for the first interface in the list by entering the following command: iscsi tpgroup create -t tag name interface -t tag is the target portal group tag you recorded for the interface before upgrading Data ONTAP. name is the name of the new target portal group. interface is the identifier of the interface for this group. Example: <code>iscsi tpgroup create -t 1 e0_old e0</code>
3	Repeat the previous step for each interface in the list.
4	Verify the new values match the old values you recorded by entering the following command again: iscsi tpgroup show Example: <pre>filer2> iscsi tpgroup show TPGTag Name Member Interfaces 1 e0_old e0 2 e9a_old e9a 3 e9b_old e9b 1000 e0_default (none) 1001 e9a_default (none) 1002 e9b_default (none)</pre>

Changing target portal group tags using FilerView: To change the target portal groups to their original values, complete the following steps:

Step	Action
1	Click LUNs > iSCSI > Manage Portal Groups. Result: A list of target portal groups is displayed.
2	Click Add Portal Group. Result: The Add iSCSI Portal Group page is displayed.
3	For the first target portal group in the list you recorded before the upgrade, enter a name for the new portal group, enter the tag value from the list, and select the interface. Example: Add iSCSI Portal Group ? Portal Group Name: e0_old ? Enter name of new portal group Portal Group Tag: 1 ? Enter value between 1 and 256 for the tag. If left blank, the system will assign the tag value Interfaces: <input type="list" value="e0"/> ? Add interfaces to this portal group. If you add interfaces to this portal group, sessions using those interfaces will be temporarily disconnected. Add
4	Click Add, and then click OK.
5	Repeat for each target portal group in the list.

Step	Action																																			
6	Verify the new values match the old values you recorded by clicking LUNs > iSCSI > Manage Portal Groups again. Example: <table><tr><th></th><th>Portal Group</th><th>Type</th><th>Tag</th><th>Interfaces</th></tr><tr><td></td><td>e0_default</td><td>system</td><td>1000</td><td></td></tr><tr><td></td><td>e9a_default</td><td>system</td><td>1001</td><td></td></tr><tr><td></td><td>e9b_default</td><td>system</td><td>1002</td><td></td></tr><tr><td>☐</td><td>e0_old</td><td>user-defined</td><td>1</td><td>e0</td></tr><tr><td>☐</td><td>e9a_old</td><td>user-defined</td><td>2</td><td>e9a</td></tr><tr><td>☐</td><td>e9b_old</td><td>user-defined</td><td>3</td><td>e9b</td></tr></table>		Portal Group	Type	Tag	Interfaces		e0_default	system	1000			e9a_default	system	1001			e9b_default	system	1002		☐	e0_old	user-defined	1	e0	☐	e9a_old	user-defined	2	e9a	☐	e9b_old	user-defined	3	e9b
	Portal Group	Type	Tag	Interfaces																																
	e0_default	system	1000																																	
	e9a_default	system	1001																																	
	e9b_default	system	1002																																	
☐	e0_old	user-defined	1	e0																																
☐	e9a_old	user-defined	2	e9a																																
☐	e9b_old	user-defined	3	e9b																																

Starting the iSCSI service

Once all of the original target portal group tags have been configured, you can restart the iSCSI service on the storage system.

Starting the iSCSI service using the command line: To start the iSCSI service on the storage system using the command line, complete the following step:

Step	Action
1	On the storage system console, enter the following command: iscsi start

Starting the iSCSI service using FilerView: To start the iSCSI service on the storage system using FilerView, complete the following steps:

Step	Action
1	Click LUNs > Enable/Disable.
2	Select the Enable check box for iSCSI.
3	Click Apply.

**For more
information**

For more information, see the *Block Access Management Guide*.

About this chapter This chapter explains how you upgrade the firmware on the storage system and on the storage system's disk array.

Topics in this chapter This chapter discusses the following topics:

- ◆ [“Upgrading the system firmware”](#) on page 60
- ◆ [“Upgrading disk firmware”](#) on page 67

Upgrading the system firmware

System firmware upgrade requirements

Make sure that you upgrade to the latest system firmware version when you upgrade Data ONTAP. In some upgrade scenarios, system firmware upgrades are mandatory.

Note

There are special system firmware requirements that must be satisfied if you are upgrading to Data ONTAP 7.1 from a release earlier than 6.5. For more information, see “[Upgrading from releases earlier than Data ONTAP 6.5](#)” on page 34.

To upgrade system firmware, you must have

- ◆ Access to the storage system’s console
- ◆ One of the following:
 - ❖ Data ONTAP 7.1 upgrade system files downloaded to your system; for CompactFlash-based storage systems and NearStore® series storage systems (F825, F880, FAS200 series, FAS900 series, FAS3000 series, R100, R150, and R200)
 - ❖ A firmware service image; for CompactFlash-based storage systems and NearStore series storage systems (F825, F880, FAS200 series, FAS900 series, FAS3000 series, R100, R150, and R200)
 - ❖ A firmware upgrade diskette; for diskette-based storage systems (F810, F820, and F840 storage systems)

Instructions for creating a firmware diskette are in Appendix A, “[Creating Diskettes](#),” on page 159.

Obtaining system firmware updates

System firmware updates for CompactFlash-based storage systems are included with distribution files for Data ONTAP upgrades. When you run the download command, the firmware service image included with the Data ONTAP upgrade is copied to your storage system’s CompactFlash card.

System firmware updates are also available at the NOW site:

- ◆ For CompactFlash-based storage systems, see Service Images on the NOW site at <http://now.netapp.com/NOW/download/tools/serviceimage/>.
- ◆ For diskette-based storage systems, see the System Firmware Release Table on the NOW site at <http://now.netapp.com/NOW/tools/releasetable.shtml>.

F87 storage systems use firmware that does not need to be upgraded.

Command for updating system firmware

The `update-flash` command loads system firmware from media (CompactFlash card or firmware diskette) into your storage system's flash Programmable ROM (PROM). The command must be run from the boot prompt—the `ok` prompt on F800 series, FAS900 series, and NearStore systems or the CFE prompt on FAS200 or FAS3000 series systems.

Determining whether your storage system needs a system firmware update

To determine whether your storage system needs a system firmware update, compare the version of installed system firmware with the latest version available on the NOW site.

To display the storage system's current system firmware, enter the following command:

```
sysconfig -a
```

The command output will include an entry similar to the following:

```
Firmware release: 4.2_i1
```

To display the most current system firmware available on the NOW site, follow the steps described in "[Storage system firmware availability](#)" on page 4.

Comparing firmware versions on CompactFlash-based systems: On CompactFlash-based systems, you can also display the system firmware version on the CompactFlash card with the `version -b` command.

You can run `version -b` immediately after you run the `download` command (which copies the new firmware to the CompactFlash card). By comparing the output of the `sysconfig -a` and `version -b` commands, you can determine if your system needs a firmware upgrade.

Examples: If the version of the newly loaded firmware displayed by the `version` command is

- ◆ The same as the installed firmware version displayed by `sysconfig`, your storage system does not need a system firmware update.
- ◆ Later than the installed firmware displayed by `sysconfig`, your storage system needs a system firmware update.

Attention

If you find that the system firmware version on the CompactFlash card is *older* than the installed version, *do not* update system firmware with the `update-flash` command.

Upgrading storage system firmware on single systems

Download and upgrade storage system firmware for your storage system according to the instructions on the NOW site.

Upgrading storage system firmware on clustered systems

If you are upgrading clustered storage systems, there are two methods available for upgrading system firmware:

- ◆ Nondisruptive (CompactFlash-based storage systems only)
The nondisruptive upgrade method is appropriate when you need to maintain service availability during system firmware upgrades.
- ◆ Standard
The standard upgrade method is appropriate when you can schedule downtime for system firmware upgrades.

For more information about selecting a clustered upgrade method, see “[Upgrade host requirements](#)” on page 6.

The following sections describe nondisruptive firmware upgrades and standard upgrades for CompactFlash-based and diskette-based systems. Be sure to refer to the firmware download and installation instructions on the NOW site for details about firmware upgrades on your platform.

Nondisruptive firmware upgrades: To upgrade system firmware using the nondisruptive upgrade process, complete the following steps.

Step	Action
1	Download the system firmware from the NOW site to your Windows or UNIX client as directed.

Step	Action	
2	On each storage system, referred to as system A and system B in the following steps, enter the following command: priv set advanced Result: The storage system prompt now displays an asterisk (*) after the storage system name to indicate that you are in advanced mode.	
3	On each storage system, enter the download -d command in priv set advanced mode as directed. If necessary, format the service partition according to the instructions.	
4	If...	Then...
	CIFS is not in use in system A	Go to Step 5 .
	CIFS is in use in system A	Enter the following command: cifs terminate -t nn where <i>nn</i> is a notification period (in seconds) appropriate for your clients. After that period of time, proceed to Step 5 .
5	At the console of system B, enter the following command: cf takeover Result: This command causes system A to shut down gracefully and leaves system B in takeover mode.	
6	To display the boot prompt (CFE or ok, depending on your storage system) at the system A console, complete one of the following procedures: ◆ Press Delete at the system A console when instructed after the boot sequence starts. ◆ When the Waiting for giveback message appears on the console of system A, press Ctrl-C at the system A console. When the Continue with boot? message appears, enter no.	

Step	Action
7	At the boot prompt, enter bye and continue installing the firmware on system A following the directions on the NOW site. Result: After completing the firmware installation on system A, you see the message <code>Waiting for giveback</code> .
8	At the console of system B, enter the following command: cf giveback Result: This command causes system A to reboot with the new system firmware and resume normal operation as a cluster partner.
9	Repeat Step 5 through Step 8 to update the partner storage system; in other words, bring down and update system B with partner A in takeover mode.

Standard firmware upgrades on CompactFlash-based systems: To upgrade system firmware using the standard upgrade process, complete the following steps.

Step	Action
1	Download the system firmware from the NOW site to your Windows or UNIX client as directed.
2	On each storage system, enter the following command: priv set advanced Result: The storage system prompt now displays an asterisk (*) after the storage system name to indicate that you are in advanced mode.
3	On each storage system, enter the <code>download -d</code> command in <code>priv set advanced</code> mode as directed. If necessary, format the service partition according to the instructions.
4	On either storage system, disable the cluster by entering the following command: cf disable

Step	Action	
5	Continue installing the firmware on each storage system following the directions on the NOW site.	
6	If a Troika card...	Then...
	Is not used as the cluster interconnect	Reenable the cluster by entering the following command on one of the storage systems: cf enable
	Is used as the cluster interconnect	Wait for the console to display the Troika firmware download complete and Resetting Troika NIC messages. Then reenale the cluster by entering the following command on one of the storage systems: cf enable

Standard firmware upgrades for diskette-based systems: To upgrade system firmware on diskette-based systems using the standard upgrade process, complete the following steps.

Step	Action
1	Download the system firmware file from the NOW site and create a diskette as directed.
2	On either storage system, disable the cluster by entering the following command: cf disable
3	Install the firmware on each node following the directions on the NOW site.

Step	Action	
4	If a Troika card...	Then...
	Is not used as the cluster interconnect	Reenable the cluster by entering the following command on one of the storage systems: cf enable
	Is used as the cluster interconnect	Wait for the console to display the Troika firmware download complete and Resetting Troika NIC messages. Then reenale the cluster by entering the following command on one of the storage systems: cf enable

Upgrading disk firmware

About disk firmware upgrades

Disk firmware is upgraded automatically when one of the following is true:

- ◆ You add new disks or a disk shelf.
- ◆ Data ONTAP detects disk firmware updates in the `/etc/disk_fw` directory after a reboot.

Each storage system is shipped with a `/etc/disk_fw` directory that contains the latest firmware revisions at that time. Disk firmware updates can be added to this directory at the following times:

- ◆ During a Data ONTAP upgrade
Disk firmware updates are often included with an upgrade to a new release family. Disk firmware updates are occasionally included in Data ONTAP upgrades within release families. This is the most common way to upgrade disk firmware.
- ◆ During a manual firmware upgrade
You might be directed to download a disk firmware upgrade from the NOW site in the event that you encounter problems with certain disk types or you receive a notice from NetApp.

Each disk drive manufacturer has its own disk drive firmware. Therefore, disk firmware upgrades can include updates to firmware for one or more disk drive types. Because your storage system might use drives from multiple drive manufacturers, whether you are affected by a disk firmware update depends on the types and numbers of drives on your system.

Note

Because disk firmware upgrades are potentially time consuming and can affect storage system availability, make sure that you verify whether a disk firmware upgrade will take place during a Data ONTAP upgrade. To do so, follow the instructions described in “[Detecting outdated disk firmware](#)” on page 71.

Service availability during disk firmware upgrades

When you upgrade to Data ONTAP 7.1, the availability of storage system services during a disk firmware upgrade depends on the type of RAID protection on aggregates containing the disks.

Disk firmware upgrades can take place in two ways:

- ◆ Background (nondisruptive) disk firmware upgrade

Nondisruptive disk firmware upgrades take place automatically in the background when the disks are members of aggregates of the following types:

- ❖ RAID-DP
- ❖ Mirrored RAID-DP (RAID-DP with SyncMirror® software)
- ❖ Mirrored RAID4 (RAID4 with SyncMirror software)

For more information, see [“About background disk firmware upgrades”](#) on page 69.

- ◆ Standard disk firmware upgrade

In Data ONTAP 7.1, disk firmware upgrades for RAID4 aggregates must complete before the new Data ONTAP version can finish booting. Storage system services are not available until the disk firmware upgrade completes.

For more information, see [“About standard disk firmware upgrades”](#) on page 70.

For example, if a storage system contains a RAID-DP and a RAID4 aggregate and disks in both aggregates require a disk firmware upgrade, the storage system cannot service requests until the RAID4 aggregate’s disk firmware is updated, even though the RAID-DP aggregate’s disks are upgrading firmware in the background.

To verify the type of RAID protection you have on your system, complete the following step.

Step	Action
1	At the storage system command line, enter the following command: <pre>aggr status</pre> Result: You see output similar to the following: <pre>Aggr State Status Options data2_vol online raid-dp, flex data1_vol online raid-dp, flex vol0 online raid4, flex root</pre> If any volume, including the root volume, is configured with RAID4 protection, a standard disk firmware upgrade (interrupting storage system services) will take place at the next reboot when new disk firmware is present on the system. Note In some storage systems, RAID4 is configured on the root volume by default. Be sure to check the RAID type of your root volume before you upgrade, and reconfigure it if necessary, if you require a nondisruptive disk firmware upgrade.

About background disk firmware upgrades: When a storage system configured with RAID-DP or SyncMirror reboots and there is new disk firmware present, the affected drives are automatically and sequentially taken offline, and the storage system responds normally to read and write requests. If any request affects an offline drive, the read requests are satisfied by reconstructing data from other disks in the RAID group, while write requests are written to a log. When the disk firmware upgrade is complete, the drive is brought back online after resynchronizing any write operations that took place while the drive was offline.

During a background disk firmware upgrade, the storage system functions normally. You will see status messages as disks are taken offline to upgrade firmware and brought back online when the firmware upgrade is complete. Background disk firmware upgrades proceed sequentially for active data disks and for spare disks. Sequential disk firmware upgrades ensure that there will be no data loss through double-disk failure.

Offline drives are marked with the annotation “offline” in the `vol status -r` command output. While a spare disk is offline, it cannot be added to a volume or selected as a replacement drive for reconstruct operations. However, a disk would normally remain offline for a very short time (a few minutes at most) and hence would not interfere with normal system operation.

The background disk firmware upgrade will be completed unless the following conditions are encountered:

- ◆ Degraded volumes on the storage system
- ◆ Disks needing a firmware update are present in a volume or plex that is in an offline state

Automatic background disk firmware upgrades will resume when these conditions are addressed. For more information about determining volume status and state, see the *Storage Management Guide*.

Automatic background disk firmware upgrades are overridden when the `disk_fw_update` command is issued. For more information about this command, see “[Command for upgrading disk firmware](#)” on page 72.

Note

Automatic background disk firmware upgrades are enabled by the `raid.background_disk_fw_update.enable` option, which is set to on by default. The value of this option can be overridden during a cluster takeover, when the `disk_fw_update` command is issued, or when a disk firmware upgrade is required for disks in a RAID4 aggregate. You are advised not to change the default value unless directed by NetApp staff.

About standard disk firmware upgrades: In Data ONTAP 7.0.1 and later RAID4 aggregates (as well as in all volume and aggregate configurations in earlier Data ONTAP releases), standard disk firmware upgrades take place automatically during the first reboot after the appearance of new disk firmware on the system. Because disk drives must be spun down and spun back up to install new firmware, disk firmware upgrades can take many minutes depending on the number of drive types and disk drives per type on your storage system. During a disk firmware upgrade, the disks of the affected drive types are not available.

Note

If you upgrade RAID protection to RAID-DP, disk firmware upgrades take place in the background and are non-disruptive.

Detecting outdated disk firmware

AutoSupport messages include information about disk firmware installed on your storage system. The NOW Installed Product pages use these messages to monitor the firmware versions on your storage system and to post notices when installed disk firmware in the /etc/disk_fw has been superseded. These notices indicate that the disk firmware on at least some of your disks will be upgraded during your next Data ONTAP upgrade, which can help you plan your upgrade.

To use the NOW Installed Products service to monitor disk firmware versions:

- ◆ AutoSupport must be enabled on your storage system.
For more information about AutoSupport, see the *System Administration Guide*.
- ◆ You must have registered your NetApp products.
For more information, see Product Registration on the NOW site at <http://now.netapp.com/Self-Service/Forms/frmProductReg.asp>.

To determine if your storage system has outdated disk firmware, complete the following steps.

Step	Action
1	Use a Web browser to go to http://now.netapp.com/ .
2	Click Installed Products.
3	Display the product details for the storage systems you are upgrading by entering search criteria for a specific system or displaying a list of systems at your company.
4	In the AutoSupport Status category, click Analysis. Result: You see a list of System Level Notices for your storage system. If there is a notice that newer disk firmware is available for your storage system, a disk firmware upgrade will take place with your next Data ONTAP upgrade. If there is no disk firmware notice, the disk firmware on your system is up to date.

When to upgrade disk firmware manually

If you receive error messages about the firmware compatibility, you must manually upgrade the disk firmware using one of the following procedures:

- ◆ Manually update disk firmware with the `disk_fw_update` command.
See “[Command for upgrading disk firmware](#)” on page 72.
- ◆ Download new disk firmware from the NOW site, then run the `disk_fw_update` command.
See “[Obtaining disk firmware updates](#)” on page 72.

Note

When you upgrade the storage system software, disk firmware is updated automatically as part of the storage system software upgrade process. A manual update is not necessary unless the new firmware is not compatible with the storage system disks.

Command for upgrading disk firmware

You use the `disk_fw_update` command from the storage system console on a running system to update firmware on all disks or a specified disk on a storage system. The `disk_fw_update` command upgrades disks for which firmware files are present in the `/etc/disk_fw` directory and which need to be updated. It does not upgrade other disks.

The `disk_fw_update` command is applicable to both SCSI disks and Fibre Channel disks.

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

Note

This command makes disks inaccessible for up to two minutes, so network sessions using the storage system should be terminated before running the command. This is particularly true for CIFS sessions, which otherwise are terminated while this command executes.

This command overrides any background disk firmware upgrade that is in progress.

Obtaining disk firmware updates

System firmware updates are available from the Disk Drive & Firmware Matrix, on the NOW site at <http://now.netapp.com/NOW/download/tools/diskfw/>. Download and upgrade disk firmware according to the instructions.

About this chapter This chapter provides you with the information you need to upgrade Data ONTAP software to the latest release.

Topics in this chapter This chapter discusses the following topics:

- ◆ [“Prerequisites to software upgrade”](#) on page 74
- ◆ [“Installing Data ONTAP system files”](#) on page 77
- ◆ [“Downloading Data ONTAP to single systems and rebooting”](#) on page 92
- ◆ [“Downloading Data ONTAP to clustered systems and rebooting”](#) on page 98
- ◆ [“Concluding the upgrade”](#) on page 135
- ◆ [“Where to go from here”](#) on page 137

Prerequisites to software upgrade

What you need for the upgrade

To upgrade the storage system software to the latest release of Data ONTAP, you need the following items:

- ◆ System files
You can get the latest system files from the following locations:
 - ❖ On the NOW site at <http://now.netapp.com/>
 - ❖ On the CD-ROM that is supplied in your media package
- ◆ The latest version of the firmware for your model of storage system
The version information is available on the NOW site at <http://now.netapp.com/NOW/tools/releasetable.shtml>.

Note

The system files package includes the latest version of system firmware that can be downloaded on CompactFlash-based systems.

F87 storage systems use firmware that does not need to be upgraded.

- ◆ Data ONTAP system boot diskettes
You need these only if your storage system is diskette-based (F810, F820, and F840 systems) and the storage system is not running at all.
For information on creating boot diskettes, see Appendix A, “[Creating Diskettes](#),” on page 159.

You can also find utilities on the NOW site that can help ensure an efficient upgrade. See the Tools & Utilities page at <http://now.netapp.com/NOW/cgi-bin/toolsutils/>.

Preparing for the upgrade

Before installing the latest Data ONTAP release on your storage system:

- ◆ Verify that you have resolved any upgrade issues identified in “[Resolving upgrade issues](#)” on page 16.
- ◆ Verify whether you need to update storage system firmware as described in “[Determining whether your storage system needs a system firmware update](#)” on page 61.
- ◆ Verify whether you need to update disk firmware as described in “[Detecting outdated disk firmware](#)” on page 71.

- ◆ If you have clustered storage systems, verify that they are correctly configured using the Cluster Configuration Checker as described in “[Before upgrading clustered storage systems](#)” on page 8.
- ◆ If you are running SnapMirror, identify storage systems with destination volumes and upgrade them before upgrading storage systems with source volumes as described in “[Planning for upgrading storage systems running SnapMirror](#)” on page 12.

You should also check whether you need to perform one or both of the procedures described in the following table.

Procedure	Complete the procedure if...
Verifying that the storage system has a domain account	You are running CIFS on the storage system and are using a Windows NT 4.0 domain controller for authentication. In this case, the storage system must have a domain account.
Enabling DNS with Windows 2000 name server addresses	You are running CIFS on the storage system and are using a Windows 2000 domain controller for authentication.

Enabling DNS with Windows 2000 name server addresses

To enable DNS with Windows 2000 name server addresses, complete the following steps.

Step	Action
1	Using a text editor, create or open the /etc/resolv.conf file in the root volume. Enter up to three lines, each specifying a Windows 2000 name server host in the following format: nameserver <i>ip_address</i> Example: nameserver 192.9.200.10
2	Save the file.

Step	Action
3	Enter the following command at the storage system console to enable DNS: options dns.enable on

Verifying that the storage system has a domain account

To verify that the storage system has a domain account, complete the following step.

Step	Action
1	From the storage system, enter the following command: cifs lookup

Installing Data ONTAP system files

About the installation of system files

The procedure for installing system files depends on the type of upgrade host from which you are installing the files. For more information, see [“About upgrade hosts”](#) on page 6.

For detailed information

For details about installing the Data ONTAP system files, see the following topics:

- ◆ [“UNIX client steps”](#) on page 78
- ◆ [“Windows client steps”](#) on page 82
- ◆ [“Using the Upgrade Utility”](#) on page 86
- ◆ [“HTTP server steps”](#) on page 87

UNIX client steps

When to perform UNIX client steps

Perform the steps described in this section only if you are using a UNIX client to install the upgrade.

Procedure overview

This part of the upgrade involves performing the following procedures in order:

1. Mounting the storage system on your client
2. Downloading the system files to your client from NOW or copying them from the CD-ROM
3. Installing the system files

Note

To perform the upgrade, you need access to both the storage system's system console and a client system.

Mounting the storage system on your client

To mount the storage system on your UNIX client, complete the following steps.

Step	Action
1	Mount the storage system's root file system to the client's /mnt directory: mount <i>filer</i>:/ /mnt <i>filer</i> is the name of the storage system. <i>/mnt</i> is the directory on the client where you want to mount the storage system's root file system.
2	Use the UNIX cd command to change to the /mnt directory: cd /mnt <i>/mnt</i> is the directory on the client where you mounted the storage system's root file system.

Step	Action	
3	If you are...	Then...
	Downloading system files from NOW	Complete the procedure described in “ Downloading the system files from NOW ” on page 79.
	Copying system files from the CD-ROM	Complete the procedure described in “ Copying the system files from the CD-ROM ” on page 80.

Downloading the system files from NOW

To download the system files from the NOW site to your client, complete the following steps.

Step	Action
1	Use a Web browser to log in to NOW.
2	Click Service & Support.
3	Click Download Software.
4	In the Software Download table, click the Select Platform list box in the Data ONTAP product row.
5	Select your storage system type from the list and click Go.
6	Follow the prompts to reach the software download page.
7	Choose the desired UNIX system files software and download it to the mount point that you chose in the previous procedure.
8	From the client, extract the files with the following command: <pre>tar xvf /mnt/filename</pre> <i>/mnt</i> is the directory on the client where you want to mount the storage system’s root file system. <i>filename</i> is the name you downloaded from NOW. Result: The tar command extracts files and reports the files’ names: tar_image.Z and install_netapp.

Step	Action
9	Complete the procedure described in “ Installing the system files ” on page 80.

Copying the system files from the CD-ROM

To access the system files from the CD-ROM and copy them to your client, complete the following steps.

Step	Action
1	Mount the Data ONTAP CD-ROM onto your client.
2	From the client, extract the files by entering the following command: <pre>tar xvf /CD-ROM_MNT_PT/ONTAP/release/type/SYSFILES.TAR</pre> <i>/CD-ROM_MNT_PT</i> is the name of your CD-ROM mount point. <i>release</i> is the release number on the CD-ROM. <i>type</i> is one of the following values: ◆ X86 for an F87 storage system, and for F800, FAS900, or NearStore series storage systems ◆ MIPS for a FAS200 or FAS3000 series storage system Result: The tar command extracts the files and reports the names; they are tar_image.Z and install_netapp.
3	Complete the procedure described in “ Installing the system files ” on page 80.

Installing the system files

To install the system files, complete the following steps.

Step	Action	
1	If you want to...	Then...
	Delete the tar files after the files are extracted	From the client, enter the following command: ./install_netapp /mnt <i>/mnt</i> is the directory on the client where you mounted the storage system's root file system.
	Keep the tar files after the files are extracted	From the client, enter the following command: ./install_netapp -k /mnt <i>/mnt</i> is the directory on the client where you mounted the storage system's root file system.
	Result: The install script decompresses and extracts the files.	
2	From the client, unmount the storage system's root file system (/) by entering the following commands: cd / umount /mnt <i>/mnt</i> is the directory on the client where you mounted the storage system's root file system.	
3	Complete the appropriate procedure to download Data ONTAP: ◆ For single storage systems, go to "Downloading Data ONTAP to single systems and rebooting" on page 92. ◆ For clustered storage systems, go to "Downloading Data ONTAP to clustered systems and rebooting" on page 98.	

Windows client steps

When to perform Windows client steps

Perform the steps described in this section only if you are using a Windows PC to install the upgrade.

Procedure overview

This part of the upgrade involves performing the following steps:

1. Mapping the storage system to a drive on your client
2. Downloading system files to your client from the NOW site or copying them from the CD-ROM and installing them

Mapping the storage system to a drive

To map the root directory of your storage system to a drive on your client, complete the following steps.

Step	Action	
1	Log in to your client as Administrator or log in using an account that has full control on the storage system C\$ directory.	
2	Map a drive to the C\$ directory of your storage system.	
3	If you are...	Then...
	Downloading system files from the NOW site	Complete the procedure described in “Downloading and installing the system files from NOW” on page 83.
	Copying system files from the CD-ROM	Complete the procedure described in “Copying and installing the system files from the CD-ROM” on page 84.

Downloading and installing the system files from NOW

To download the system files from the NOW site and install them onto your client, complete the following steps.

Step	Action
1	Use a Web browser to log in to NOW.
2	Click Service & Support.
3	Click Download Software.
4	In the Software Download table, click the Select Platform list box in the Data ONTAP product row.
5	Select your storage system type from the list and click Go.
6	Follow the prompts to reach the software download page.
7	Choose the desired Windows system files software for your storage system and download it to the drive to which you mapped the storage system in the previous procedure.
8	Go to the drive to which you downloaded the software.
9	Double-click the file that you downloaded. Result: A dialog box appears.
10	In the WinZip dialog box, enter the letter of the drive to which you mapped the storage system. For example, if you chose drive N, replace <DRIVE>:\ETC with the following path: N:\ETC
11	Ensure that the following check boxes are selected: ◆ Overwrite Files Without Prompting ◆ When Done Unzipping Open... Leave the other options as they are.
12	Click the Unzip button. Result: In the dialog box's lower panel, confirmation messages are displayed as files are decompressed.

Step	Action
13	Complete the appropriate procedure to download Data ONTAP: ◆ For single storage systems, go to “Downloading Data ONTAP to single systems and rebooting” on page 92. ◆ For clustered storage systems, go to “Downloading Data ONTAP to clustered systems and rebooting” on page 98.

Copying and installing the system files from the CD-ROM

To access the system files from the Data ONTAP CD-ROM and copy them to your client, complete the following steps.

Step	Action
1	Insert the Data ONTAP CD-ROM into your CD-ROM drive.
2	Go to one of the following folders on the Data ONTAP CD-ROM: ◆ ONTAP\release\X86 for an F87 storage system, and for F800, FAS900, or NearStore series storage systems ◆ ONTAP\release\MIPS for a FAS200 or FAS 3000 series storage system <i>release</i> is the release number on the CD-ROM.
3	Double-click setup.exe. Result: The setup program displays a dialog box.
4	Make sure that you followed the steps described in the dialog box, then click OK. Result: The WinZip dialog box appears.

Step	Action
5	Follow the WinZip prompts to unzip the file that you downloaded onto the drive to which you mapped the storage system. When appropriate, ensure that the following check boxes are selected: ◆ Overwrite Files Without Prompting ◆ When Done Unzipping Open... Leave the other options as they are. Result: In the dialog box's lower panel, confirmation messages are displayed as files are decompressed.
6	Complete the appropriate procedure to download Data ONTAP: ◆ For single storage systems, go to “Downloading Data ONTAP to single systems and rebooting” on page 92. ◆ For clustered storage systems, go to “Downloading Data ONTAP to clustered systems and rebooting” on page 98.

Installing Data ONTAP system files

Using the Upgrade Utility

About the Upgrade Utility

The Data ONTAP Upgrade Utility is a Perl script that allows you to upgrade storage systems from UNIX or Windows clients. The script is also available in executable form for Windows clients. The utility is especially useful if you do not have Common Internet File System (CIFS) or Network File System (NFS) licensed—for example, on a storage system serving only Fiber Channel Protocol (FCP)—and the conventional methods for copying system files to the storage system are not available.

For detailed information

For more information about acquiring and using the Upgrade Utility, see the NOW site at http://now.netapp.com/NOW/download/tools/upgrade_ontap/.

Note

You must be running Java on the storage system to use the Upgrade Utility.

HTTP server steps

About using an HTTP server to download Data ONTAP

If you have an HTTP server with the latest Data ONTAP upgrades accessible to your storage system, you can use the `software` command to download and install Data ONTAP to your storage system. When you use HTTP to download Data ONTAP, you do not have to mount the storage system to a UNIX administration host or map a drive to the storage system using Windows to perform the installation.

You can download Data ONTAP upgrades to both single and clustered storage systems.

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

Note

Transport for `software` command transactions uses the HTTP protocol, not HTTPS. As such, it is not a secure transport method.

Preparing the HTTP server

To prepare the HTTP server, copy the `setup.exe` file (for example, `70_setup_i.exe`) from the NOW site or another system to the directory on the HTTP server from which you will serve the file.

Copying Data ONTAP to your storage system without installing

You can copy a Data ONTAP upgrade to your storage system without immediately installing it, for instance, if you want to perform the installation at a later time. The downloaded file is saved on the storage system in the `/etc/software` directory.

To copy Data ONTAP onto your storage system without installing the software, complete the following step from the storage system console.

Step	Action
1	Enter the following command: <pre>software get url -f filename</pre> <i>url</i> is the HTTP location from which you want to download Data ONTAP. Use the following URL syntax if you need to specify a username, password, host or port to access files on the HTTP server using Basic Access Authentication (RFC2617): <pre>http://username:password@host:port/path</pre> Use the -f flag to overwrite an existing software file of the same name in the storage system's <code>/etc/software</code> directory. If a file of the same name exists and you do not use the -f flag, the download will fail and you will be prompted to use -f. <i>filename</i> is the file name you specify for the software file being downloaded to your storage system. If no destination file name is specified, Data ONTAP uses the file name listed in the URL from which you are downloading and places the copy in the <code>/etc/software</code> directory on the storage system. Example: The following is an example of the <code>software get</code> command using a new destination file name. <pre>software get http://www.myserver.com/netapp/R70/pc/70_setup_i.exe 70_setup_i.exe</pre> Result: You see a message similar to the following: <pre>software: copying to /etc/software/70_setup_i.exe software: 100% file read from location. software: /etc/software/70_setup_i.exe has been copied.</pre>

Installing Data ONTAP using HTTP

You can use the `software install` command to copy and install the Data ONTAP upgrade files to your storage system simultaneously.

To install Data ONTAP from an HTTP server, complete the following steps from the storage system console.

Step	Action	
1	Enter the following command: <pre>software install url</pre> <i>url</i> is the HTTP location from which you want to download Data ONTAP. Use the following URL syntax if you need to specify a username, password, host or port to access files on the HTTP server using Basic Access Authentication (RFC2617): <pre>http://username:password@host:port/path</pre> Result: The software is downloaded and installed on your storage system, and you see a message similar to the following: <pre>tpubs-cf2> software install http://www.myserver.com/netapp/R70/pc/70_setup_i.exe software: installing software, this could take a few minutes... software: installation completed. Please type download to load the new software and reboot subsequently for changes to take effect. tpubs-cf2></pre>	
2	If you are upgrading...	Then...
	A single storage system	Go to “ Downloading Data ONTAP to single systems and rebooting ” on page 92.
	Clustered storage systems	Go to “ Downloading Data ONTAP to clustered systems and rebooting ” on page 98.

Installing a file that has already been downloaded

If you previously downloaded Data ONTAP using the `software get` command, you can install the software by completing the following steps from the storage system console.

Step	Action	
1	Enter the following command: software install <i>software_filename</i> <i>software_filename</i> is the name of the file you downloaded using the <code>software get</code> command. Result: The software is downloaded and installed on your storage system, and you see a message similar to the following: <pre>tpubs-cf2> software install 70_setup_i.exe software: installing software, this could take a few minutes... software: installation completed. Please type download to load the new software and reboot subsequently for changes to take effect. tpubs-cf2></pre>	
2	If you are upgrading...	Then...
	A single storage system	Go to “Downloading Data ONTAP to single systems and rebooting” on page 92.
	Clustered storage systems	Go to “Downloading Data ONTAP to clustered systems and rebooting” on page 98.

Listing software files in the /etc/software directory

To list the files that have been downloaded to the `/etc/software` directory, complete the following step from the storage system console.

Step	Action
1	Enter the following command: software list

Deleting software files in the /etc/software directory

To delete files in the /etc/software directory, complete the following step from the storage system console.

Step	Action
1	To delete one or more files, enter the following command: software delete <i>software_file1</i> <i>software_file2</i> <i>software_file1</i> and <i>software_file2</i> are the names of the files you want to delete.

Downloading Data ONTAP to single systems and rebooting

Selecting a single system upgrade process

To upgrade your single system, follow the instructions in the following sections based on your storage system model:

- ◆ [“Downloading Data ONTAP and system firmware on an FAS200 series or FAS3000 series storage system”](#) on page 92
- ◆ [“Downloading Data ONTAP and system firmware on an F825 or F880 storage system, or a FAS900 or NearStore series storage system”](#) on page 95
- ◆ [“Downloading Data ONTAP on F87, F810, F820, and F840 storage systems”](#) on page 97

Downloading Data ONTAP and system firmware on an FAS200 series or FAS3000 series storage system

The FAS200 series and FAS3000 series storage systems come with a CompactFlash card that you use when downloading firmware and system software.

Before initiating this download procedure, verify that you have completed any prerequisite procedures described in [“Preparing for the upgrade”](#) on page 74.

To download the Data ONTAP software and, if necessary, upgrade the firmware on one of these storage systems, complete the following steps at the storage system console.

Step	Action
1	Enter the following command to copy the kernel and firmware data files to the CompactFlash card: download Result: The download command provides a status message similar to the following: <pre>download: Depending on system load, it may take many download: minutes to complete this operation. Until it download: finishes, you will not be able to use the download: console. You can cancel this operation by download: hitting Ctrl-C in the next 6 seconds.</pre> <pre>Tue Jul 15 10:03:22 GMT [download.request:notice]: Operator requested download initiated download: Downloading boot device..... Mon Apr 21 15:30:24 PDT [nbns_timer:info]: All CIFS name registrations complete for local server.... download: Downloading boot device (Service Area)</pre> Then the following storage system message appears: <pre>Tues Jul 15 10:11:51 GMT [download.requestDone:notice]: Operator requested download complete</pre>
2	Enter the following command to shut down the storage system: halt Result: After the storage system shuts down, the CFE firmware prompt appears. Note_____ For more information about the Common Firmware Environment (CFE) boot environment and commands, see the <i>System Administration Guide</i>. _____

Step	Action	
3	If a firmware upgrade...	Then...
	Is needed on your storage system	Go to Step 4 .
	Is not needed on your storage system	Go to Step 6 .
4	At the CFE prompt, enter the following command: update_flash Result: The update_flash command provides an acknowledgement similar to the following. Reading flash0a: Done. 493328 bytes read ** DO NOT TURN OFF YOUR MACHINE UNTIL THE FLASH UPDATE COMPLETES!! ** Programming...done. 493328 bytes written Reading fatfs://ide0.0/MIPS/firmware/FAS250/firmware.img: Done. 493392 bytes read Flash image is 493328 bytes, flags 00000002, CRC DE227E0D ** DO NOT TURN OFF YOUR MACHINE UNTIL THE FLASH UPDATE COMPLETES!! ** Programming...done. 493328 bytes written	
5	Enter y to continue. Result: The system updates the firmware, displays several status messages, and displays the CFE prompt.	
6	At the CFE prompt, enter the following command to reboot the system using the new software and, if applicable, the new firmware: bye	

Downloading Data ONTAP and system firmware on an F825 or F880 storage system, or a FAS900 or NearStore series storage system

The F825 and F880 storage systems, FAS900 series storage systems, and R100, R150, and R200 NearStore systems use a CompactFlash card for firmware and system software downloads.

Before initiating this download procedure, verify that you have completed any prerequisite procedures described in “[Preparing for the upgrade](#)” on page 74.

To download the Data ONTAP software and, if necessary, upgrade the firmware on one of these storage systems, complete the following steps at the storage system console.

Step	Action
1	Enter the following command to copy the executables to the boot blocks of the storage system’s hard disks: download Result: The download command provides a status message similar to the following: <pre>Tue Jun 19 10:03:22 GMT [download.request:notice]: Operator requested download initiated download: Downloading boot device..... download: Downloading disks: opening Data ONTAP 7.1 done download: boot blocks successfully copied to disk(s)</pre> Then a message similar to the following appears: <pre>Tues Jun 19 10:11:51 GMT [download.requestDone:notice]: Operator requested download completed</pre>
2	Enter the following command to shut down the storage system: halt Result: After the storage system shuts down, the ok prompt appears.
3	Enter the following command to reboot the system: bye Result: The storage system begins the boot process.

Step	Action	
4	If a firmware upgrade...	Then...
	Is needed on your storage system	When the memory test begins, press Del to interrupt the boot process. Note If you are using a terminal emulation package and it does not support Del, use Ctrl-Bksp instead. Result: The console displays the following messages: <pre> Skipping auto-boot Complete Startup aborted </pre> Then the console displays the ok prompt. Continue to Step 5.
	Is not needed on your storage system	You are done.
5	Enter the following command: update-flash Result: The system prompts you to continue.	
6	Enter y to continue. Result: The system updates the firmware, displays several status messages, and displays the ok prompt. Note If you installed a system firmware upgrade, you might have to run the <code>set-defaults</code> command from the firmware ok prompt <i>before</i> rebooting your storage system. For more information, see “Considerations concerning system firmware upgrades” on page 35.	

Step	Action
7	Enter the following command to reboot the system using the new firmware and software: bye

Downloading Data ONTAP on F87, F810, F820, and F840 storage systems

Before initiating this download procedure, verify that you have completed any prerequisite procedures described in “[Preparing for the upgrade](#)” on page 74.

To download Data ONTAP to the disks and reboot F87, F810, F820, and F840 storage systems, complete the following steps.

Step	Action
1	At the storage system console, enter the following command to copy the executables to the boot blocks of the storage system’s hard disks: download Result: The download command provides an acknowledgment similar to the following: <pre>Tue Jun 19 10:03:22 GMT [download.request:notice]: Operator requested download initiated download: Downloading disks: opening Data ONTAP 7.1 done download: boot blocks successfully copied to disk(s)</pre> Then the following storage system prompt appears: <pre>Tues Jun 19 10:11:51 GMT [download.requestDone:notice]: Operator requested download completed</pre>
2	At the storage system console, reboot the storage system with the following command: reboot
3	As needed, perform the procedures described in “Concluding the upgrade” on page 135.

Downloading Data ONTAP to clustered systems and rebooting

About this section

Select the proper upgrade method as described in “[Choosing a clustered upgrade procedure](#)” on page 99.

Then upgrade your clustered storage systems using one of the following methods:

- ◆ “[Nondisruptive upgrades from the Data ONTAP 7.0 release family](#)” on page 101
- ◆ “[Nondisruptive upgrades from the Data ONTAP 6.5 release family](#)” on page 109
- ◆ “[Nondisruptive upgrades within release families](#)” on page 117
- ◆ “[Standard upgrades](#)” on page 124

Attention

Do not enter the Data ONTAP download command on a clustered storage system that is in takeover mode.

Choosing a clustered upgrade procedure

Which procedure to use for clustered system upgrades

The procedure you use to download Data ONTAP to clustered storage systems depends on the storage system type and the kind of upgrade—nondisruptive or standard.

- ◆ Nondisruptive upgrades are most appropriate for the following circumstances:
 - ❖ Your storage systems use CompactFlash cards (F825 and F880 storage systems, and FAS270, FAS900, and FAS3000 series storage systems).
 - ❖ You are upgrading to the Data ONTAP 7.1 release family from Data ONTAP 6.5.3 or later, or you are upgrading within the Data ONTAP 7.1 release family (for example, from Data ONTAP 7.1 to 7.1P1).
 - ❖ You are using RAID-DP or there is no disk firmware upgrade associated with the software upgrade.
 - ❖ High availability of storage system services is critical.

For instructions on performing nondisruptive upgrades from the Data ONTAP 7.0 release family (for example, from Data ONTAP 7.0.0.1 to 7.1), see [“Nondisruptive upgrades from the Data ONTAP 7.0 release family”](#) on page 101.

For instructions on performing nondisruptive upgrades from the Data ONTAP 6.5 release family (for example, from Data ONTAP 6.5.4 to 7.1), see [“Nondisruptive upgrades from the Data ONTAP 6.5 release family”](#) on page 109.

For instructions on performing nondisruptive upgrades within the Data ONTAP 7.1 release family (for example, from Data ONTAP 7.1 to 7.1P1), see [“Nondisruptive upgrades within release families”](#) on page 117.

Note

If you are planning to perform a nondisruptive upgrade on a system with FCP or iSCSI hosts, see [“About nondisruptive cluster upgrades”](#) on page 8.

- ◆ Standard upgrades are most appropriate when you can schedule downtime for major software changes.

You must use standard upgrade procedures when you upgrade to Data ONTAP 7.1 from releases earlier than Data ONTAP 6.5.3

For instructions on performing standard upgrades, see the section appropriate to your storage system type:

- ❖ [“Downloading Data ONTAP to clustered FAS270 storage systems and FAS3000 series storage systems”](#) on page 124
- ❖ [“Downloading Data ONTAP to clustered F810, F820, and F840 storage systems”](#) on page 132
- ❖ [“Downloading Data ONTAP to clustered F810, F820, and F840 storage systems”](#) on page 132

For more information about which upgrade procedure is most appropriate for your storage systems and upgrade circumstances, see [“Planning for upgrading clusters”](#) on page 8.

For more information about determining whether system or disk firmware upgrades are required, see:

- ◆ [“Determining whether your storage system needs a system firmware update”](#) on page 61
- ◆ [“Detecting outdated disk firmware”](#) on page 71

Nondisruptive upgrades from the Data ONTAP 7.0 release family

Upgrading your clustered systems using the nondisruptive method

Before initiating the nondisruptive upgrade procedure, verify that you have completed any prerequisite procedures described in [“Preparing for the upgrade”](#) on page 74. All CompactFlash-based systems that satisfy these prerequisites can use the nondisruptive upgrade procedure.

To download Data ONTAP to clustered storage systems for a nondisruptive upgrade to Data ONTAP 7.1 from the 7.0 release family, complete the following steps.

Step	Action
1	At the console of each storage system, enter the following command to copy the executables to the boot blocks of the storage system’s hard disks: download Result: The download command provides an acknowledgment similar to the following: <pre>Tue Jun 19 10:03:22 GMT [download.request:notice]: Operator requested download initiated download: Downloading boot device..... download: Downloading disks: opening Data ONTAP 7.1 done download: boot blocks successfully copied to disk(s)</pre> Then a message similar to the following appears: <pre>Tues Jun 19 10:11:51 GMT [download.requestDone:notice]: Operator requested download completed</pre> Note_____ The storage system console is unavailable until the download procedure is complete. _____

Step	Action
2	At the console of each storage system, enter the following command to verify that clustering is enabled: cf status Result: The <code>cf status</code> command output should be similar to the following: Cluster enabled, filerA is up. If it is not, enter the following command to enable clustering: cf enable Then verify that clustering is re-enabled by entering the <code>cf status</code> command.
3	At the console of each storage system, enter the following commands to compare the installed version of system firmware with the version on the CompactFlash card: sysconfig -a version -b Result: The <code>sysconfig</code> command output contains an entry similar to the following: Firmware release: 4.3_i1 The <code>version</code> command output contains an entry similar to the following: 1:/x86/firmware/xfiler/firmware.img: Firmware 4.3.1_i2 If the version of the newly loaded firmware displayed by the <code>version</code> command is ◆ The same as the installed firmware version displayed by <code>sysconfig</code>, your storage system does not need a system firmware update. ◆ Later than the installed firmware displayed by <code>sysconfig</code>, your storage system needs a system firmware update. ◆ Earlier than the installed firmware displayed by <code>sysconfig</code>, <i>do not</i> update system firmware with the <code>update-flash</code> command.

Step	Action	
4	If...	Then...
	CIFS is not in use in system A	Go to Step 5 .
	CIFS is in use in system A	Enter the following command: cifs terminate -t nn where <i>nn</i> is a notification period (in minutes) appropriate for your clients. After that period of time, proceed to Step 5 .
5	At the console of system B, enter the following command: cf takeover Result: This command causes system A to shut down gracefully and leaves system B in takeover mode. Note_____ When system A displays the waiting for giveback message, you can also halt and power off system A to perform maintenance procedures. For more information, see “About nondisruptive cluster upgrades” on page 8.	
6	If...	Then...
	You do not need to upgrade system firmware	Go to Step 13 .
	If you need to update system firmware	Go to Step 7 .

Step	Action
7	To display the boot prompt (CFE or ok, depending on your storage system) at the system A console, complete one of the following procedures: ◆ Press Del at the system A console when instructed after the boot sequence starts. ◆ When the waiting for giveback message appears on the console of system A, press Ctrl-C at the system A console. When the Continue with boot? message appears, enter no.
8	At the boot prompt, enter the following command: bye
9	During the memory test, press Del to interrupt the boot process. Note_____ If you have a terminal emulation package and it does not support Del, use Ctrl-Bksp instead. Result: The console displays the following messages: Skipping auto-boot Complete Startup aborted Then the console displays the boot prompt (CFE or ok).
10	Enter the following command: update-flash Result: The system prompts you to continue.

Step	Action
11	Enter y to continue. Result: The system updates the firmware, displays several status messages, and displays the boot prompt. Note_____ If you installed a system firmware upgrade on a storage system running Open Firmware (OFW), you might have to run the <code>set - defaults</code> command from the firmware ok prompt <i>before</i> rebooting your storage system. For more information, see “Considerations concerning system firmware upgrades” on page 35.
12	Enter the following command to reboot the system using the new firmware and software: bye
13	When the <code>Waiting for giveback</code> message appears on the console of system A, enter the following command at the console of system B: cf giveback -f Result: This command causes system A to reboot with the new system configuration—a Data ONTAP version and any system firmware and hardware changes—and resume normal operation as a cluster partner. Note_____ At this point in the upgrade procedure—system A is running Data ONTAP 7.1 and system B is running an earlier Data ONTAP release family—the systems are in a state of “version mismatch.” This means that normal clustering functions such as NVRAM mirroring and automatic takeover are not in effect, and error messages indicating version mismatch and mailbox format problems are seen. This is expected behavior that is temporary and not harmful. Nonetheless, you should complete the upgrade procedure as quickly as possible; do not allow the two systems to remain in a state of version mismatch longer than necessary.

Step	Action	
14	If...	Then...
	CIFS is not in use in system B	Go to Step 15 .
	CIFS is in use in system B	Enter the following command: cifs terminate -t nn where <i>nn</i> is a notification period (in minutes) appropriate for your clients. After that period of time, proceed to Step 15 .
15	At the console of system A, enter the following command: cf takeover -n Result: You see output similar to the following: Waiting for partner to be cleanly shutdown using the 'halt' command. Press Ctrl-C to abort wait...	
16	At the console of system B, enter the following command: halt Result: This command causes system B to shut down cleanly, flushing file-system information in memory to disk. Note When system B displays the Waiting for giveback message, you can also power off system B to perform maintenance procedures. For more information, see “About nondisruptive cluster upgrades” on page 8.	
17	If...	Then...
	You do not need to upgrade system firmware	Go to Step 23 .
	If you need to update system firmware	Go to Step 18 .

Step	Action
18	To display the boot prompt (CFE or ok, depending on your storage system) at the system B console, complete one of the following procedures: ◆ Press Del at the system B console when instructed after the boot sequence starts. ◆ When the waiting for giveback message appears on the console of system B, press Ctrl-C at the system B console. When the Continue with boot? message appears, enter no.
19	At the boot prompt, enter the following command: bye
20	During the memory test, press Del to interrupt the boot process. Note_____ If you have a terminal emulation package and it does not support Del, use Ctrl-Bksp instead. Result: The console displays the following messages: Skipping auto-boot Complete Startup aborted Then the console displays the boot prompt (CFE or ok).
21	Enter the following command: update-flash Result: The system prompts you to continue.

Step	Action
22	Enter y to continue. Result: The system updates the firmware, displays several status messages, and displays the boot prompt. Note_____ If you installed a system firmware upgrade on a storage system running Open Firmware (OFW), you might have to run the <code>set - defaults</code> command from the firmware ok prompt <i>before</i> rebooting your storage system. For more information, see “Considerations concerning system firmware upgrades” on page 35.
23	At the console of system B, enter the following command to reboot the system using the new system firmware (if it was installed) and software: bye
24	When the <code>Waiting for giveback</code> message appears on the console of system B, enter the following command at the console of system A: cf giveback -f Result: This command causes system B to reboot with the new system configuration—a Data ONTAP version and any system firmware and hardware changes—and resume normal operation as a cluster partner. When the reboot completes, the two cluster nodes are running the same Data ONTAP version.
25	As needed, perform the procedures described in “Concluding the upgrade” on page 135.

Nondisruptive upgrades from the Data ONTAP 6.5 release family

Upgrading your clustered systems using the nondisruptive method

Before initiating the nondisruptive upgrade procedure, verify that you have completed any prerequisite procedures described in “[Preparing for the upgrade](#)” on page 74. All CompactFlash-based systems that satisfy these prerequisites can use the nondisruptive upgrade procedure.

To download Data ONTAP to clustered storage systems for a nondisruptive upgrade to Data ONTAP 7.1 from the Data ONTAP 6.5 release family (Data ONTAP 6.5.3 or later), complete the following steps.

Step	Action
1	At the console of each storage system, enter the following command to copy the executables to the boot blocks of the storage system’s hard disks: download Result: The download command provides an acknowledgment similar to the following: <pre>Tue Jun 19 10:03:22 GMT [download.request:notice]: Operator requested download initiated download: Downloading boot device..... download: Downloading disks: opening Data ONTAP 7.1 done download: boot blocks successfully copied to disk(s)</pre> Then a message similar to the following appears: <pre>Tues Jun 19 10:11:51 GMT [download.requestDone:notice]: Operator requested download completed</pre> Note_____ The storage system console is unavailable until the download procedure is complete. _____

Step	Action
2	At the console of each storage system, enter the following command to verify that clustering is enabled: cf status Result: The <code>cf status</code> command output should be similar to the following: Cluster enabled, filerA is up. If it is not, enter the following command to enable clustering: cf enable Then verify that clustering is re-enabled by entering the <code>cf status</code> command.
3	At the console of each storage system, enter the following commands to compare the installed version of system firmware with the version on the CompactFlash card: sysconfig -a version -b Result: The <code>sysconfig</code> command output contains an entry similar to the following: Firmware release: 4.3_i1 The <code>version</code> command output contains an entry similar to the following: 1:/x86/firmware/xfiler/firmware.img: Firmware 4.3.1_i2 If the version of the newly loaded firmware displayed by the <code>version</code> command is ◆ The same as the installed firmware version displayed by <code>sysconfig</code>, your storage system does not need a system firmware update. ◆ Later than the installed firmware displayed by <code>sysconfig</code>, your storage system needs a system firmware update. ◆ Earlier than the installed firmware displayed by <code>sysconfig</code>, <i>do not</i> update system firmware with the <code>update-flash</code> command.

Step	Action	
4	If...	Then...
	CIFS is not in use in system A	Go to Step 5 .
	CIFS is in use in system A	Enter the following command: cifs terminate -t nn where <i>nn</i> is a notification period (in minutes) appropriate for your clients. After that period of time, proceed to Step 5 .
5	At the console of system B, enter the following command: cf takeover Result: This command causes system A to shut down gracefully and leaves system B in takeover mode. Note_____ When system A displays the waiting for giveback message, you can also halt and power off system A to perform maintenance procedures. For more information, see “About nondisruptive cluster upgrades” on page 8.	
6	If...	Then...
	You do not need to upgrade system firmware	Go to Step 13 .
	If you need to update system firmware	Go to Step 7 .

Step	Action
7	To display the boot prompt (CFE or ok, depending on your storage system) at the system A console, complete one of the following procedures: ◆ Press Del at the system A console when instructed after the boot sequence starts. ◆ When the waiting for giveback message appears on the console of system A, press Ctrl-C at the system A console. When the Continue with boot? message appears, enter no.
8	At the boot prompt, enter the following command: bye
9	During the memory test, press Del to interrupt the boot process. Note— If you have a terminal emulation package and it does not support Del, use Ctrl-Bksp instead. Result: The console displays the following messages: Skipping auto-boot Complete Startup aborted Then the console displays the boot prompt (CFE or ok).
10	Enter the following command: update-flash Result: The system prompts you to continue.

Step	Action
11	Enter y to continue. Result: The system updates the firmware, displays several status messages, and displays the boot prompt. Note_____ If you installed a system firmware upgrade on a storage system running Open Firmware (OFW), you might have to run the <code>set - defaults</code> command from the firmware <code>ok</code> prompt <i>before</i> rebooting your storage system. For more information, see “Considerations concerning system firmware upgrades” on page 35.
12	Enter the following command to reboot the system using the new firmware and software: bye
13	When the <code>Waiting for giveback</code> message appears on the console of system A, enter the following command at the console of system B: cf giveback -f Result: This command causes system A to reboot with the new system configuration—a Data ONTAP version and any system firmware and hardware changes—and resume normal operation as a cluster partner. Note_____ At this point in the upgrade procedure—system A is running Data ONTAP 7.1 and system B is running an earlier Data ONTAP release family—the systems are in a state of “version mismatch.” This means that normal clustering functions such as NVRAM mirroring and automatic takeover are not in effect, and error messages indicating version mismatch and mailbox format problems are seen. This is expected behavior that is temporary and not harmful. Nonetheless, you should complete the upgrade procedure as quickly as possible; do not allow the two systems to remain in a state of version mismatch longer than necessary.

Step	Action	
14	If...	Then...
	CIFS is not in use in system B	Go to Step 15 .
	CIFS is in use in system B	Enter the following command: cifs terminate -t nn where <i>nn</i> is a notification period (in minutes) appropriate for your clients. After that period of time, proceed to Step 15 .
15	At the console of system B, enter the following command: halt Result: This command causes system B to shut down cleanly, flushing file-system information in memory to disk.	
16	At the console of system A, enter the following command: cf takeover -n Result: You see output similar to the following: <pre>The state of the partner cannot be determined. The partner must be halted. Do you want to continue (y/n)?</pre> Having halted system B, enter y to continue. Note_____ When system B displays the <code>Waiting for giveback</code> message, you can also power off system B to perform maintenance procedures. For more information, see “About nondisruptive cluster upgrades” on page 8. _____	

Step	Action	
17	If...	Then...
	You do not need to upgrade system firmware	Go to Step 23 .
	If you need to update system firmware	Go to Step 18 .
18	To display the boot prompt (CFE or ok, depending on your storage system) at the system B console, complete one of the following procedures: ◆ Press Del at the system B console when instructed after the boot sequence starts. ◆ When the Waiting for giveback message appears on the console of system B, press Ctrl-C at the system B console. When the Continue with boot? message appears, enter no.	
19	At the boot prompt, enter the following command: bye	
20	During the memory test, press Del to interrupt the boot process. Note If you have a terminal emulation package and it does not support Del, use Ctrl-Bksp instead. Result: The console displays the following messages: Skipping auto-boot Complete Startup aborted Then the console displays the boot prompt (CFE or ok).	
21	Enter the following command: update-flash Result: The system prompts you to continue.	

Step	Action
22	Enter y to continue. Result: The system updates the firmware, displays several status messages, and displays the boot prompt. Note_____ If you installed a system firmware upgrade on a storage system running Open Firmware (OFW), you might have to run the <code>set - defaults</code> command from the firmware ok prompt <i>before</i> rebooting your storage system. For more information, see “Considerations concerning system firmware upgrades” on page 35.
23	At the console of system B, enter the following command to reboot the system using the new system firmware (if it was installed) and software: bye
24	When the <code>Waiting for giveback</code> message appears on the console of system B, enter the following command at the console of system A: cf giveback -f Result: This command causes system B to reboot with the new system configuration—a Data ONTAP version and any system firmware and hardware changes—and resume normal operation as a cluster partner. When the reboot completes, the two cluster nodes are running the same Data ONTAP version.
25	As needed, perform the procedures described in “Concluding the upgrade” on page 135.

Downloading Data ONTAP to clustered systems and rebooting

Nondisruptive upgrades within release families

Upgrading your clustered systems using the nondisruptive method

Before initiating the nondisruptive upgrade procedure, verify that you have completed any prerequisite procedures described in “[Preparing for the upgrade](#)” on page 74. All CompactFlash-based systems that satisfy these prerequisites can use the nondisruptive upgrade procedure.

To download Data ONTAP to clustered storage systems for a nondisruptive upgrade within a release family (for example, from Data ONTAP 7.1 to 7.1.1), complete the following steps at the console of one of the storage systems in the cluster.

Step	Action
1	At the console of each storage system, enter the following command to copy the executables to the boot blocks of the storage system's hard disks: download Result: The download command provides an acknowledgment similar to the following: <pre>Tue Jun 19 10:03:22 GMT [download.request:notice]: Operator requested download initiated download: Downloading boot device..... download: Downloading disks: opening Data ONTAP 7.1 done download: boot blocks successfully copied to disk(s)</pre> Then a message similar to the following appears: <pre>Tues Jun 19 10:11:51 GMT [download.requestDone:notice]: Operator requested download completed</pre> Note_____ The storage system console is unavailable until the download procedure is complete. _____

Step	Action
2	At the console of each storage system, enter the following command to verify that clustering is enabled: cf status Result: The <code>cf status</code> command output should be similar to the following: Cluster enabled, filerA is up. VIA Interconnect is up (link 0 up, link 1 down). If the output says that clustering is not enabled, enter the following command to enable clustering: cf enable
3	At the console of each storage system, enter the following commands to compare the installed version of system firmware with the version on the CompactFlash card: sysconfig -a version -b Result: The <code>sysconfig</code> command output contains an entry similar to the following: Firmware release: 4.3_i1 The <code>version</code> command output contains an entry similar to the following: 1:/x86/firmware/xfiler/firmware.img: Firmware 4.3.1_i2 If the version of the newly loaded firmware displayed by the <code>version</code> command is ◆ The same as the installed firmware version displayed by <code>sysconfig</code>, your storage system does not need a system firmware update ◆ Later than the installed firmware displayed by <code>sysconfig</code>, your storage system needs a system firmware update ◆ Earlier than the installed firmware displayed by <code>sysconfig</code>, <i>do not</i> update system firmware with the <code>update-flash</code> command

Step	Action	
4	If...	Then...
	CIFS is not in use in system A	Go to Step 5 .
	CIFS is in use in system A	Enter the following command: cifs terminate -t nn where <i>nn</i> is a notification period (in minutes) appropriate for your clients. After that period of time, proceed to Step 5 .
5	At the console of system B, enter the following command: cf takeover Result: This command causes system A to shut down gracefully and leaves system B in takeover mode. Note_____ When system A displays the <code>Waiting for giveback</code> message, you can also halt system A to perform other upgrade and maintenance procedures. For more information, see “About nondisruptive cluster upgrades” on page 8. _____	
6	If...	Then...
	You do not need to upgrade system firmware	Go to Step 13 .
	If you need to update system firmware	Go to Step 7 .

Step	Action
7	To display the boot prompt (CFE or ok, depending on your storage system) at the system A console, complete one of the following procedures: ◆ Press Del at the system A console when instructed after the boot sequence starts. ◆ When the waiting for giveback message appears on the console of system A, press Ctrl-C at the system A console. When the Continue with boot? message appears, enter no.
8	At the boot prompt, enter the following command: bye
9	During the memory test, press Del to interrupt the boot process. Note_____ If you have a terminal emulation package and it does not support Del, use Ctrl-Bksp instead. Result: The console displays the following messages: Skipping auto-boot Complete Startup aborted Then the console displays the boot prompt (CFE or ok).
10	Enter the following command: update-flash Result: The system prompts you to continue.

Step	Action
11	Enter y to continue. Result: The system updates the firmware, displays several status messages, and displays the boot prompt. Note_____ If you installed a system firmware upgrade on a storage system running Open Firmware (OFW), you might have to run the set - defaults command from the firmware ok prompt <i>before</i> rebooting your storage system. For more information, see “Considerations concerning system firmware upgrades” on page 35.
12	Enter the following command to reboot the system using the new firmware and software: bye

Step	Action	
13	If...	Then...
	FCP or iSCSI is not in use in system A	When the waiting for giveback message appears on the console of system A, enter the following command at the console of system B: cf giveback Result: This command causes system A to reboot with the new system configuration—a Data ONTAP version or other system firmware and hardware changes—and resume normal operation as a cluster partner.
	FCP or iSCSI is in use in system A	When the waiting for giveback message appears on the console of system A, wait for at least 8 minutes to allow host multipathing software to stabilize. Then enter the following command at the console of system B: cf giveback Result: This command causes system A to reboot with the new system configuration—a Data ONTAP version or other system firmware and hardware changes—and resume normal operation as a cluster partner.

Step	Action	
14	If...	Then...
	FCP or iSCSI is not in use in system A	Repeat Step 4 through Step 13 to update the partner storage system; in other words, bring down and update system B with partner A in takeover mode.
	FCP or iSCSI is in use in system A	After system A resumes normal operation as a cluster partner, wait for at least 8 minutes to allow host multipathing software to stabilize. Then repeat Step 4 through Step 13 to update the partner storage system; in other words, bring down and update system B with partner A in takeover mode.
15	As needed, perform the procedures described in “ Concluding the upgrade ” on page 135.	

Standard upgrades

Selecting a standard clustered system upgrade process

To upgrade your clustered systems using the standard method, follow the instructions in one of the following sections based on your storage system model:

- ◆ [“Downloading Data ONTAP to clustered FAS270 storage systems and FAS3000 series storage systems”](#) on page 124
- ◆ [“Downloading Data ONTAP to clustered F825 and F880 storage systems, and clustered FAS900 series storage systems”](#) on page 128
- ◆ [“Downloading Data ONTAP to clustered F810, F820, and F840 storage systems”](#) on page 132

Downloading Data ONTAP to clustered FAS270 storage systems and FAS3000 series storage systems

FAS270 storage systems and FAS3000 series storage systems use a CompactFlash card for firmware and system software downloads.

Before initiating the standard upgrade procedure, verify that you have completed any prerequisite procedures described in [“Preparing for the upgrade”](#) on page 74.

To download Data ONTAP to clustered FAS270 storage systems or FAS3000 series storage systems using the standard upgrade procedure, complete the following steps at the console of one of the storage systems in the cluster.

Step	Action
1	Disable the cluster by entering the following command at the console of one of the storage systems: cf disable

Step	Action
2	Enter the following command to copy the executables to the boot blocks of the storage system's hard disks: download Result: The download command provides a status message similar to the following: <pre>download: Depending on system load, it may take many download: minutes to complete this operation. Until it download: finishes, you will not be able to use the download: console. You can cancel this operation by download: hitting Ctrl-C in the next 6 seconds.</pre> <pre>Tue Jul 15 10:03:22 GMT [download.request:notice]: Operator requested download initiated download: Downloading boot device..... Mon Apr 21 15:30:24 PDT [nbns_timer:info]: All CIFS name registrations complete for local server.... download: Downloading boot device (Service Area)</pre> Then the following storage system message appears: <pre>Tues Jul 15 10:11:51 GMT [download.requestDone:notice]: Operator requested download complete</pre>

Step	Action
3	Enter the following commands to compare the installed version of system firmware with the version on the CompactFlash card: <pre>sysconfig -a version -b</pre> Result: The sysconfig command output contains an entry similar to the following: <pre>Firmware release: CFE 1.1.0</pre> The version command output contains an entry similar to the following: <pre>1:/mips/firmware/fas270/firmware.img: Firmware 1.1.0</pre> If the version of the newly loaded firmware displayed by the version command is ◆ The same as the installed firmware version displayed by sysconfig, your storage system does not need a system firmware update. ◆ Later than the installed firmware displayed by sysconfig, your storage system needs a system firmware update. ◆ Earlier than the installed firmware displayed by sysconfig, <i>do not</i> update system firmware with the update-flash command.
4	Enter the following command to shut down the storage system: <pre>halt</pre> Result: After the storage system shuts down, the CFE firmware prompt appears. Note_____ For more information about the Common Firmware Environment (CFE) boot environment and commands, see the section on booting the FAS200 and FAS3000 series storage systems in the <i>System Administration Guide</i>. _____

Step	Action	
5	If a firmware upgrade...	Then...
	Is needed on your storage system	Go to Step 6 .
	Is not needed on your storage system	Go to Step 8 .
6	At the CFE prompt, enter the following command: update_flash Result: The update_flash command provides an acknowledgement similar to the following. Reading flash0a: Done. 493328 bytes read ** DO NOT TURN OFF YOUR MACHINE UNTIL THE FLASH UPDATE COMPLETES!! ** Programming...done. 493328 bytes written Reading fatfs://ide0.0/MIPS/firmware/FAS270/firmware.img: Done. 493392 bytes read Flash image is 493328 bytes, flags 00000002, CRC DE227E0D ** DO NOT TURN OFF YOUR MACHINE UNTIL THE FLASH UPDATE COMPLETES!! ** Programming...done. 493328 bytes written	
7	Enter y to continue. Result: The system updates the firmware, displays several status messages, and displays the CFE prompt.	
8	At the CFE prompt, enter the following command to reboot the system using the new software and, if applicable, the new firmware: bye	

Step	Action
9	While the cluster is disabled, repeat Step 4 through Step 8 at the storage system console of the partner storage system. Attention — Do not proceed to Step 10 until both storage systems in the cluster have been rebooted with the new version of Data ONTAP.
10	Reenable the cluster by entering the following command on one of the storage systems: cf enable
11	Update the disk firmware by entering the following command first at the console of one of the clustered storage systems and then at the console of the partner storage system: disk_fw_update
12	As needed, perform the procedures described in “ Concluding the upgrade ” on page 135.

Downloading Data ONTAP to clustered F825 and F880 storage systems, and clustered FAS900 series storage systems

The F825 and F880 storage systems and FAS900 series storage systems use a CompactFlash card for firmware and system software downloads.

Before initiating the standard upgrade procedure, verify that you have completed any prerequisite procedures described in “[Preparing for the upgrade](#)” on page 74.

To download Data ONTAP to an F825 and F880 cluster, or a FAS900 series cluster using the standard upgrade procedure, complete the following steps at the storage system console.

Step	Action
1	On either storage system, disable the cluster by entering the following command: cf disable

Step	Action
2	Enter the following command to copy the executables to the boot blocks of the storage system's hard disks: download Result: The download command provides an acknowledgment similar to the following: <pre>Tue Jun 19 10:03:22 GMT [download.request:notice]: Operator requested download initiated download: Downloading boot device..... download: Downloading disks: opening Data ONTAP 7.1 done download: boot blocks successfully copied to disk(s)</pre> Then a message similar to the following appears: <pre>Tues Jun 19 10:11:51 GMT [download.requestDone:notice]: Operator requested download completed</pre>

Step	Action
3	Enter the following commands to compare the installed version of system firmware with the version on the CompactFlash card: <pre>sysconfig -a version -b</pre> Result: The sysconfig command output contains an entry similar to the following: <pre>Firmware release: 4.3_i1</pre> The version command output contains an entry similar to the following: <pre>1:/x86/firmware/xfiler/firmware.img: Firmware 4.3.1_i2</pre> If the version of the newly loaded firmware displayed by the version command is ◆ The same as the installed firmware version displayed by sysconfig, your storage system does not need a system firmware update. ◆ Later than the installed firmware displayed by sysconfig, your storage system needs a system firmware update. ◆ Earlier than the installed firmware displayed by sysconfig, <i>do not</i> update system firmware with the update-flash command.
4	Enter the following command to shut down the storage system: <pre>halt</pre> Result: After the storage system shuts down, the ok prompt appears.
5	Enter the following command: <pre>bye</pre>

Step	Action	
6	During the memory test, press Del to interrupt the boot process. Note If you have a terminal emulation package and it does not support Del, use Ctrl-Bksp instead. Result: The console displays the following messages: Skipping auto-boot Complete Startup aborted Then the console displays the ok prompt.	
7	If a firmware upgrade...	Then...
	Is needed on your storage system	Go to Step 8 .
	Is not needed on your storage system	Go to Step 10 .
8	Enter the following command: update-flash Result: The system prompts you to continue.	
9	Enter y to continue. Result: The system updates the firmware, displays several status messages, and displays the ok prompt. Note If you installed a system firmware upgrade, you might have to run the set-defaults command from the firmware ok prompt <i>before</i> rebooting your storage system. For more information, see “Considerations concerning system firmware upgrades” on page 35.	
10	Enter the following command to reboot the system using the new firmware and software: bye	

Step	Action	
11	While the cluster is disabled, repeat Step 2 through Step 10 at the storage system console of the partner storage system. Attention Do not proceed to Step 13 until both storage systems in the cluster have been rebooted with the new version of Data ONTAP.	
12	If a Troika card...	Then...
	Is not used as the cluster interconnect	Reenable the cluster by entering the following command on one of the storage systems: cf enable
	Is used as the cluster interconnect	Wait for the console to display the Troika firmware download complete and Resetting Troika NIC messages. Then reenble the cluster by entering the following command on one of the storage systems: cf enable
13	Update the disk firmware by entering the following command first at the console of one of the clustered storage systems and then at the console of the partner storage system: disk_fw_update	
14	As needed, perform the procedures described in “Concluding the upgrade” on page 135.	

Downloading Data ONTAP to clustered F810, F820, and F840 storage systems

Before initiating the standard upgrade procedure, verify that you have completed any prerequisite procedures described in “[Preparing for the upgrade](#)” on page 74.

To download Data ONTAP to clustered F810, F820, or F840 storage systems using the standard upgrade procedure, complete the following steps at the console of one of the storage systems in the cluster.

Step	Action
1	Disable the cluster by entering the following command: cf disable
2	Enter the following command to copy the executables to the boot blocks of the storage system's hard disks: download Result: The download command provides a status message similar to the following: <pre>Tue Jun 19 10:03:22 GMT [download.request:notice]: Operator requested download initiated download: Downloading disks: opening Data ONTAP 7.1 done download: boot blocks successfully copied to disk(s)</pre> Then a message similar to the following appears: <pre>Tues Jun 19 10:11:51 GMT [download.requestDone:notice]: Operator requested download completed</pre>
3	Reboot the storage system with the following command: reboot
4	While the cluster is disabled, repeat Step 2 and Step 3 on the partner storage system. Attention _____ Do not proceed to Step 5 until both storage systems in the cluster have been rebooted with the new version of Data ONTAP. _____

Step	Action	
5	If a Troika card...	Then...
	Is not used as the cluster interconnect	Reenable the cluster by entering the following command on one of the storage systems: cf enable
	Is used as the cluster interconnect	Wait for the console to display the Troika firmware download complete and Resetting Troika NIC messages. Then reenale the cluster by entering the following command on one of the storage systems: cf enable
6	a. Update the disk firmware by entering the following command at the console of one of the clustered storage systems: disk_fw_update b. Repeat Step 6a on the partner storage system.	
7	As needed, perform the procedures described in “Concluding the upgrade” on page 135.	

Concluding the upgrade

Check what you need to do to conclude

After reviewing and addressing issues identified in Chapter 2, “[Resolving upgrade issues](#),” on page 16, check the following table to determine whether you need to perform additional procedures to conclude your upgrade.

Procedure	Complete the procedure if...
Entering codes for new licenses	You purchased any new service or protocol that you are starting to use with the upgrade.
Creating a serial number file	You want to have a serial number file. A serial number file makes it easier for technical support to assist you after receiving an AutoSupport e-mail message from your storage system.

Entering codes for new licenses

To enter new license codes, complete the following step from the storage system console.

Step	Action
1	Enter the following command: <pre>license add code ...</pre> <i>code</i> is the license code that you were given when you purchased the license for a service or protocol. Example: In this example, the license code for CIFS is ABCDEFG. <pre>license add CIFS ABCDEFG</pre> Result: The storage system acknowledges each license, as shown in the following example: <pre>A cifs site license has been installed. Use "cifs setup" to configure cifs. Run cifs setup to enable cifs.</pre>

Creating a serial number file

If the `/vol/vol0/etc/serialnum` file does not exist on your storage system, complete the following steps.

Step	Action
1	Record your storage system's serial number. You can determine the serial number by entering the <code>sysconfig</code> command at the storage system console.
2	From a UNIX or Windows client, use a text editor to create a file called <code>serialnum</code> in <code>/vol/vol0/etc/</code> , and put the storage system's serial number in it.

Where to go from here

Product documentation

Product documentation for the storage system is available online on the NOW site at <http://now.netapp.com/>, and you can also order printed copies from this Web site. See the Data ONTAP 7.1 *Release Notes* for information about this Data ONTAP release.

The following documents describe Data ONTAP features and administration procedures:

- ◆ *System Administration Guide*
- ◆ *Network Management Guide*
- ◆ *Storage Management Guide*
- ◆ *Cluster Installation and Administration Guide*
- ◆ *Data Protection: Online Backup and Recovery Guide*
- ◆ *Data Protection: Tape Backup and Recovery Guide*
- ◆ *File Access and Protocol Management Guide*
- ◆ *Block Access Management Guide*
- ◆ *MultiStore® Management Guide*

Other resource documents

See the *Documentation Roadmap* for an overview of NetApp documentation resources.

About this chapter This chapter provides you with the information you need to revert Data ONTAP software to an earlier release.

Topics in this chapter This chapter discusses the following topics:

- ◆ [“Reverting to a previous release”](#) on page 140
- ◆ [“Reversion issues between release families”](#) on page 141

Reverting to a previous release

About reverting to a previous release

If you need to revert to a previous release of Data ONTAP, contact technical support.

Telephone	FAX	Email
For U.S. customers: +1 (888) 463-8277 For international customers: +1 (408) 822-6000 (ask for technical support)	+1 (408) 822-4501	support@netapp.com

Note

In some cases, you cannot revert to an earlier version of Data ONTAP. For more information, see “[Reversion issues between release families](#)” on page 141.

Reversion issues between release families

Why there might be issues

You might encounter issues if you upgrade and then decide to revert to a previous version of Data ONTAP, because features introduced in a new release might be incompatible with features of the previous release. This is especially true if you are reverting to a release earlier than the immediately previous Data ONTAP release family.

For example, if you are reverting to a release in the 6.4 family from Data ONTAP 7.0, you must review and resolve reversion issues associated with the 6.4 and 6.5 release families (but not 6.3 or earlier) before reverting from 7.0.

For more information, see Appendix C, “[Understanding Release Family Upgrades](#),” on page 171.

The following sections discuss revert issues associated with release families:

- ◆ “[Reverting to the Data ONTAP 7.0 release family](#)” on page 143
- ◆ “[Reverting to the Data ONTAP 6.5 release family](#)” on page 145
- ◆ “[Reverting to the Data ONTAP 6.4 release family](#)” on page 152
- ◆ “[Reverting to Data ONTAP 6.3 or earlier](#)” on page 157

If you need to revert to a previous release of Data ONTAP, call technical support for assistance. For more information, see “[Reverting to a previous release](#)” on page 140.

General guidelines for reverting from Data ONTAP 7.1

The following guidelines apply when reverting from Data ONTAP 7.1 to an earlier version:

- ◆ You must disable any 7.1 features before reverting.
- ◆ You cannot revert to a release earlier than Data ONTAP 6.5.
To revert the storage system to an earlier release of Data ONTAP, you must first revert to a Data ONTAP 6.5 family release and once on that release revert to an earlier version.
- ◆ You cannot revert if an upgrade is in progress.
- ◆ If you are reverting Data ONTAP on storage systems that are running the SnapMirror software, you must revert the storage systems that have SnapMirror source volumes *before* you revert the storage systems that have SnapMirror destination volumes.

This requirement applies to both asynchronous and synchronous SnapMirror for volume replication. It does not apply to SnapMirror for qtree replication.

Note

Before reverting the source storage system, you must also disable any features not supported in the earlier release. This means that after reverting, you will no longer be able to mirror certain volumes or their contents to the destination storage system, even if the destination storage system supports that feature.

For example, you must disable FlexVol volumes, as described in “[FlexVol volumes and data cannot be reverted](#)” on page 146, before reverting a SnapMirror source system to a release earlier than Data ONTAP 7.0. Although mirrored FlexVol volumes continue to exist on the destination system after the source system has been reverted, you cannot resynchronize the mirrored FlexVol volume on the destination system.

For more information, see “[Planning for upgrading storage systems running SnapMirror](#)” on page 12.

**Upgrade must
complete before
any reversion is
attempted**

If you begin the upgrade procedure to Data ONTAP 7.1 and then decide that you want to revert to an earlier Data ONTAP release, you must complete the 7.1 upgrade before reverting.

Reverting to the Data ONTAP 7.0 release family

Issues you must resolve when reverting

The following sections discuss upgrade issues that you must resolve when reverting to a Data ONTAP 7.0 family release:

- ◆ [“SnapVault relationships to SnapLock volumes must be removed before reverting”](#) on page 143
- ◆ [“Secondary SnapVault configurations must be removed from NearStore systems before reverting”](#) on page 143
- ◆ [“Reversion cannot take place during I2P initialization”](#) on page 144

SnapVault relationships to SnapLock volumes must be removed before reverting

Before reverting to a version of Data ONTAP earlier than 7.1, SnapVault relationships to SnapLock volumes must be removed. This includes

- ◆ Destroying all SnapVault relationships to WORM volumes
- ◆ Deleting all snapshots on the WORM volumes
- ◆ Destroying all volumes which have been used as LockVault™ log volumes

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

Secondary SnapVault configurations must be removed from NearStore systems before reverting

If you revert a NearStore system that is being used as a SnapVault secondary storage system to a Data ONTAP release earlier than 7.1, you must remove SnapVault configuration from all qtrees on all volumes. This includes

- ◆ SnapVault secondary qtrees
- ◆ Open Systems SnapVault secondary qtrees
- ◆ Qtrees that are the destination of VERITAS NetBackup backups (both Project 2 and Project 3 volumes)

To do so, complete the following steps at the NearStore system command line.

Step	Action
1	Identify qtrees under SnapVault control by entering the following command: snapvault status

Step	Action
2	For each qtree identified in the previous, enter the following command: snapvault stop -f /vol/my_volume/my_qtree
3	If you are managing NearStore systems with VERITAS NetBackup, enter the following commands for each volume containing NetBackup backups: snapvault stop -f /vol/my_volume/.nbu_image_data snap delete -a -f my_volume

For more information about managing secondary SnapVault relationships on NearStore systems, see the *Data Protection Online Backup and Recovery Guide*.

Reversion cannot take place during I2P initialization

When you upgrade to Data ONTAP 7.1, inode to pathname (I2P) information is written to each existing volume. You cannot revert to an earlier release while the system is being scanned to initialize I2P. Before proceeding with the reversion, you must do one of the following:

- ◆ Wait until the I2P scan completes. This can take many hours, depending on the number of inodes and volumes on your system.
- ◆ Abort the I2P scan.

To abort the I2P scan, complete the following step.

Step	Action
1	Enter the following command: vol options no_i2p on Result: Inode to pathname translation is disabled immediately and the reversion can proceed. Note _____ Because any I2P information is removed during a reversion, reverting to an earlier release can take several hours, depending on the number of inodes and volumes on your system, and whether the I2P scan completed. _____

Reverting to the Data ONTAP 6.5 release family

Issues you must resolve when reverting

The following sections discuss upgrade issues that you must resolve when reverting to a Data ONTAP 6.5 family release:

- ◆ [“Flexible root volumes cannot be reverted”](#) on page 145
- ◆ [“FlexVol volumes and data cannot be reverted”](#) on page 146
- ◆ [“FlexCache volumes cannot be reverted”](#) on page 146
- ◆ [“Systems with DiskShelf14mk2 ATshelves cannot be reverted”](#) on page 146
- ◆ [“Full administrative access privileges after reverting”](#) on page 146
- ◆ [“Changes to vol status and aggr status command output after revert”](#) on page 146
- ◆ [“SSL must be re-enabled after reverting”](#) on page 147
- ◆ [“SnapMover vFiler migration and software-based disk ownership must be disabled before reverting”](#) on page 147
- ◆ [“Connections requiring SMB signing not available after reversion”](#) on page 151
- ◆ [“Connections requiring LDAP signing not available after reversion”](#) on page 151
- ◆ [“Storage system domain accounts must be recreated after reverting to Data ONTAP 6.5.2 or earlier”](#) on page 151

Flexible root volumes cannot be reverted

If you created a flexible root volume in Data ONTAP 7.0 or later, you must convert it to a traditional root volume and destroy the FlexVol volume and any containing aggregate before reverting to an earlier release. For information about changing the root volume to a traditional volume, see the *System Administration Guide*.

Note

The reversion procedure will not allow you to proceed unless there is a traditional root volume on the system.

FlexVol volumes and data cannot be reverted

If you created and populated FlexVol volumes in Data ONTAP 7.0 or later, you must migrate the data to traditional volumes, and destroy the FlexVol volumes, any clones of FlexVol volumes, and any containing aggregates before you revert to an earlier release. For information about migrating data between volume types, see the *Storage Management Guide*.

Note

The reversion procedure will not allow you to proceed while there are FlexVol volumes, including clones, on the system.

FlexCache volumes cannot be reverted

If you created a FlexCache volume in Data ONTAP 7.0.1 or later, you must destroy it before reverting to an earlier release. For information about FlexCache volumes, see the *Storage Management Guide*.

Systems with DiskShelf14mk2 ATshelves cannot be reverted

If you have a storage system connected to DiskShelf14mk2 AT shelves with AT-FCX modules installed, you cannot revert the system to any release earlier than Data ONTAP 7.0.1. The shelves require Data ONTAP 7.0.1 or later to be run on the attached storage systems.

Systems affected: This reversion restriction applies to the following storage systems with DiskShelf14Mk2 AT shelves attached:

- ◆ FAS3050
- ◆ FAS3020
- ◆ FAS960

Full administrative access privileges after reverting

If you created user accounts in Data ONTAP 7.0 or later with groups, roles, and capabilities, their specific privilege levels will not be preserved if you revert to an earlier release. The user accounts and names will be preserved during the reversion, but each account will have full privileges in the earlier release.

Changes to vol status and aggr status command output after revert

The `vol status` command has been modified in the Data ONTAP 7.0 release family to distinguish between flexible and traditional volumes. If you created scripts for Data ONTAP 7.0 or later that rely on `vol status` command output, the scripts might not work when you revert to an earlier release. Specifically, the

Status field in Data ONTAP 7.0 or later includes comma-separated entries that indicate volume type, but in earlier releases the entries are not comma-separated and do not contain volume type information.

Similarly, if you created scripts for Data ONTAP 7.0 or later that rely on `aggr status` command output, these scripts must be rewritten to support pre-7.0 `vol status` command output.

For more information about the `vol status` command and output, see the `vol(1)` man page.

SSL must be re-enabled after reverting

If you installed and licensed the SecureAdmin product with an earlier Data ONTAP release, your existing SSL (Secure Socket Layer) protocol certificates are converted when you upgrade to Data ONTAP 7.0 or later and copies of the old certificates are preserved.

If you then revert to a release of Data ONTAP earlier than 7.0, you can continue to use the old SSL certificates as before; the new certificates are no longer valid after reversion. However, you must re-enable SSL after the reversion for the old certificates to take effect. For more information about SSL and SecureAdmin, see your SecureAdmin documentation.

Note

If you want to continue using SecureAdmin with SSL after reverting to an earlier Data ONTAP release, you must have the SecureAdmin product installed and licensed on your system. For more information about SecureAdmin, see the *SecureAdmin Administrator's Guide*.

SnapMover vFiler migration and software-based disk ownership must be disabled before reverting

Data ONTAP 6.5.1 introduced software-based disk ownership for clustered systems with DS14 or DS14Mk2 disk shelves attached and with SnapMover® licenses enabled. If you want to revert to a version of Data ONTAP earlier than 6.5.1 on a system where SnapMover vFiler migration is enabled and software-based disk ownership is assigned, you must ensure that the hardware-based ownership rules based on physical A loop and B loop connections are satisfied before reverting.

Attention

If hardware-based ownership rules are not satisfied, the storage systems might not be able to boot and their data might be unavailable after reverting.

To disable SnapMover vFiler migration and the software-based disk ownership that enables it, complete the following steps.

Step	Action
1	On one of the cluster nodes, enter the following command to ascertain whether or not the current software-based disk ownership assignments align with the cluster's A loop/B loop topology: aggr status -r Results: In the output of this command, all disks assigned to the current cluster node are listed as volume disks or as spare disks; all disks assigned to the cluster's partner are listed as partner disks. All disks on the same shelf must be assigned to the same cluster node and all disk shelf assignments must be consistent with the A loop B loop topology of the cluster. For example, in the CHAN column of the command output, all disks owned by the local storage system should have a CHAN value of FC:A and all partner disks should have a CHAN value of FC:B.
2	If the aggr status command output indicates that disk ownership reflects the A loop B loop topology, go to Step 4. If the aggr status command output indicates that disk ownership deviates from the A loop B loop topology, you must change the disk assignments back to hardware-based ownership as directed in Step 3.

Step	Action	
3	If...	Then...
	The disk ownership deviates from the A loop B loop topology because you have migrated a vFiler unit with the <code>vfiler migrate -m nocopy</code> command...	Change it back with the same command. Then go to Step 4 .
	The disk ownership deviates from the A loop B loop topology because you have split ownership of a single expansion disk shelf between the two cluster nodes...	Use the <code>disk assign</code> command to assign all disks on the shelf to the cluster node attached to the shelf's A loop connection. Then go to Step 4 .
	The disk ownership deviates from the A loop B loop topology because you have reassigned a spare disk with the <code>disk assign</code> command...	Change it back with the <code>disk assign</code> command. Then go to Step 4 .
4	Disable clustering by entering the following command. cf disable	
5	On each cluster node, uninstall the license for SnapMover. license delete snapmover	
6	Reboot each cluster node, and when prompted, press Ctrl-C to display the boot menu options.	
7	On each cluster node, enter the choice for booting in maintenance mode.	

Step	Action
8	In maintenance mode on each cluster node, enter the following command: <pre>disk remove_ownership</pre> This command disables software-based disk ownership. Note _____ When you revert to an earlier release in the Data ONTAP 6.5 release family, you must explicitly erase software-based ownership with this command. When you revert to a Data ONTAP release earlier than 6.5, software-based disk ownership is erased automatically. However, you are advised to completely disable SnapMover and software-based ownership before initiating the reversion process.
9	Halt each cluster node to exit maintenance mode: <pre>halt</pre>
10	Reboot each cluster node in normal mode. The system returns to ownership based on A loop and B loop connections. You can now safely revert the systems or enter the <code>cf enable</code> command to re-enable cluster operation.

For more information about disabling SnapMover and disk assignments, see the *MultiStore Management Guide*. For more information about erasing software-based disk ownership, see the *Storage Management Guide*.

Systems affected: Reversion rules for software-based disk ownership apply to the following clustered storage systems with SnapMover and MultiStore licenses, and with DS14 or DS14Mk2 disk shelves attached:

- ◆ FAS900 series
- ◆ F880
- ◆ F825

These rules do not apply to clustered FAS270 and FAS 3000 series systems, which also include DS14Mk2 disk shelves and which must have software-based disk ownership enabled. FAS270 systems cannot be reverted to a version of Data ONTAP earlier than 6.5, and FAS3000 series systems cannot be reverted to a version of Data ONTAP earlier than 7.0.1.

Connections requiring SMB signing not available after reversion

SMB signing support was introduced in Data ONTAP 7.0.1. If any clients have been configured to require SMB signing (security signatures) and you revert your storage system to an earlier Data ONTAP release, the clients will refuse to connect to the storage system after reverting.

In addition, if your storage system authenticates with a Windows Domain Controller that requires SMB signing and you revert to an earlier Data ONTAP release, you cannot connect to the Domain Controller after reverting.

For more information about SMB signing, see the *File Access and Protocols Management Guide*.

Connections requiring LDAP signing not available after reversion

LDAP signing support was introduced in Data ONTAP 7.0.1. If your storage system communicates with a Windows Domain Controller that requires LDAP signing and you revert to an earlier Data ONTAP release, you cannot connect to the Domain Controller after reverting.

For more information about LDAP signing, see the *File Access and Protocols Management Guide*.

Storage system domain accounts must be recreated after reverting to Data ONTAP 6.5.2 or earlier

Data ONTAP 6.5.3 and later releases correct a problem with storage system domain accounts, the Active Directory domain accounts created for your storage system. If you create a storage system domain account and then revert to a Data ONTAP release earlier than 6.5.3, you must run `cifs setup` after the reversion to recreate the storage system domain account.

Note

If you created a storage system domain account before upgrading to Data ONTAP 6.5.3 or later, it will continue to work after the upgrade, and it will also continue to work if you revert to an earlier release.

For more information about creating storage system domain accounts, see the *Software Setup Guide* for the Data ONTAP release to which you are reverting.

Reverting to the Data ONTAP 6.4 release family

Issues you must resolve before reverting

The following sections discuss upgrade issues that you must resolve before upgrading to the Data ONTAP 6.4 release family:

- ◆ [“Reversion with larger inodes”](#) on page 152
- ◆ [“Excessive subdirectories and hard links”](#) on page 154
- ◆ [“Large directories can be reverted but cannot be made larger”](#) on page 154
- ◆ [“Reverting entries in the /etc/exports file”](#) on page 154
- ◆ [“RAID-related errors”](#) on page 155
- ◆ [“Running a RAID label version that is too new”](#) on page 155
- ◆ [“SnapLock volumes must be removed before reverting”](#) on page 156
- ◆ [“Fractional reserve must be 100% before reverting”](#) on page 156

Reversion with larger inodes

When you upgrade to Data ONTAP 6.5 or later, your storage system is upgraded to support larger inodes. There are two instances in which larger inodes can cause a reversion to an earlier release to fail:

- ◆ While the inode upgrade is still in progress
- ◆ When there is not enough disk space or there are not enough free inodes in a volume

Note

Reverting to an earlier release can take several hours, depending on the number of total inodes and inodes used on your system.

To determine whether an inode upgrade is still in progress, complete the following steps.

Step	Action
1	Enter the following command: <pre>priv set advanced</pre> Result: The storage system prompt now displays an asterisk (*) after the storage system name to indicate that you are in advanced mode.

Step	Action
2	Enter the following command: <pre>waf1 scan status</pre> Example: If the upgrade is still in progress, the storage system displays a message similar to the following: <pre>Volume vol1: Scan id Type of scan progress 2 inofile upgrade 849387 of 33554472</pre> Note_____ If an inofile upgrade scan type is not displayed, the inode upgrade is done. _____

When more disk space or free inodes are needed for reverting, you see a message similar to the following:

```
Volume vol1: There are not enough free inodes to revert the ACLs
with nlinks over 65535.
A minimum of ### inodes needed (maximum ###).
```

Minimum and maximum inodes needed depends on the number of hard links (nlinks) over 65,535.

You can free up inodes in a volume by deleting files or by moving files to a volume that has free inodes.

To view the number of used and free inodes, complete the following step.

Step	Action
1	At the storage system console, enter the following command: <pre>df -i</pre>

Excessive subdirectories and hard links

If you have more than 65,533 subdirectories or 65,535 hard links, you must move the extra subdirectories to another directory and delete the extra hard links before you revert. When too many hard links exist during a revert, you will see an error message similar to the following:

```
Scanning all inodes in volume "vol1" for non-ACL inodes with nlinks
over 65535 0% done
inode 102, file has too high of an nlink (81519)
inode 101, file has too high of an nlink (90001)
0% done.
```

Large directories can be reverted but cannot be made larger

Data ONTAP 6.5 and later releases allow larger directories than in previous versions. If you have a directory larger than 65,535 KB, the directory will be accessible when you revert to an earlier version of Data ONTAP and will remain at its current size, but you cannot make it any larger.

Reverting entries in the /etc/exports file

When you upgrade to Data ONTAP 6.5 or later, new access rules are applied to the /etc/exports file. If you revert to an earlier version of Data ONTAP, the /etc/exports file will be scanned to determine whether it conforms to the earlier version's format. If not, the reversion will not proceed until the file has been fixed.

To fix the file, you can edit it manually or use the `exportfs -d` command to revert the file. Using `exportfs -d` will fix most of the revert issues. However, some changes to the /etc/exports file are not fixed with the `exportfs -d` command and cause a message to be sent to the console during a reversion. If any of these issues are not fixed, that line in the /etc/exports file does not load and access is denied for that resource. This happens in the following circumstances:

- ◆ Lines greater than 4,096 characters, excluding comments
- ◆ Entries exceeding 256 characters in either the `-rw` or `-root` options
- ◆ Too many entries in the `-access` option

Attention

Data ONTAP has no way to determine whether entries in an access list represent a netgroup or a host name; therefore, no checking is done on names during reversion. If netgroups exist in the `-rw` and `-root` options, you must manually edit the file to conform to the earlier format; otherwise, the export line containing the netgroup does not load and access is denied for that resource.

For more information about the /etc/exports file, see the `exports(5)` man page.

RAID-related errors Updates made to RAID in Data ONTAP 6.5 or later can sometimes cause problems when you revert to an earlier version, and the reversion may fail. The following is a list of possible problems and the resulting error messages that are sent to the console:

RAID reversion issue	Message with required action
RAID-DP volumes are present	There is a raid_dp volume present which cannot be reverted. Downgrade these volumes with 'vol options <vol> raidtype raid 4'.
RAID label assimilation in progress	RAID label assimilation in progress. Wait for assimilation to complete before reverting.
Failed Disk Registry in process of being updated	Please wait one moment while RAID updates the Failed Disk Registry.
Sanitized disk present	There are sanitized disks present. Either remove them or issue a 'disk sanitize release' to return sanitized disks to the spare pool.
Disk in process of being sanitized is present	There are sanitized disks present. Either wait for the sanitize to complete, or issue a 'disk sanitize abort' to return sanitized disks to the spare pool.
Sparecore disk present	There is a sparecore disk present which cannot be reverted. Wait for the sparecore to complete before running revert.

Running a RAID label version that is too new When you install a version of Data ONTAP earlier than 6.5, you might get a message that the RAID label version you are running is either invalid or too new. This can happen when the storage system contains disks that were previously

upgraded with Data ONTAP 6.5 or later, and were not reverted using the `revert_to` command. In cases where a disk has not been properly reverted, the following actions occur:

- ◆ The disk is not assimilated.
- ◆ The disk is moved to the broken disk pool with “bad label” given as the reason.
- ◆ An AutoSupport notice is sent to NetApp and the customer.

SnapLock volumes must be removed before reverting

Before reverting to a version of Data ONTAP earlier than 6.5, SnapLock volumes must be removed. For more information about SnapLock, see the *Storage Management Guide*.

Fractional reserve must be 100% before reverting

Fractional reserve (setting space reservation to less than 100 percent) is available in Data ONTAP 6.5.1 or later, but it is not available in the 6.4 release family. If you want to revert from Data ONTAP 6.5.1 or higher to a Data ONTAP 6.4 family release and you are using fractional reserve, you must make sure you have enough available space for 100 percent overwrite reserve. If you do not have enough space when you revert, Data ONTAP displays the following message: You have an over committed volume. You are required to set the `fractional_reserve` to 100. This can be done by either disabling space reservations on all objects in the volume or making more space available for full reservations or deleting all the snapshots in the volume.

To avoid this problem when you revert to a Data ONTAP 6.4 family release, make sure that you have enough available space for 100 percent overwrite reserve. If you do not have enough space or you have received this message, make space by deleting snapshots.

Reverting to Data ONTAP 6.3 or earlier

Issues you must resolve before reverting

The following sections discuss upgrade issues that you must resolve before reverting to the Data ONTAP 6.3 release family or earlier releases:

- ◆ “[Stale file handles on vFiler units](#)” on page 157
- ◆ “[Finding reversion issues for releases earlier than Data ONTAP 6.3](#)” on page 157
- ◆ “[Reversion with home directory paths in the cifs_homedir.cfg file](#)” on page 157

Stale file handles on vFiler units

Data ONTAP 6.4 and later releases include a change in the Network File System (NFS) file handle format that is not backward compatible with vFiler systems. If you have defined vFiler storage on your system, and you upgrade to 6.4 or later and then revert to an earlier version, NFS clients might report stale file handles on the vFiler units. To prevent this problem, unmount any NFS clients that are mounted on vFiler units. You can remount them after reverting.

Finding reversion issues for releases earlier than Data ONTAP 6.3

To find reversion issue information for releases earlier than Data ONTAP 6.3, go to the Information Library for Archived Documentation on the NOW site, at http://now.netapp.com/NOW/knowledge/docs/archive_docs.shtml, for a list of documentation for earlier releases of Data ONTAP.

Reversion with home directory paths in the cifs_homedir.cfg file

If you upgrade to Data ONTAP 6.3 or later and convert home directory path entries from the `cifs.home_dir` option to the `cifs_homedir.cfg` file, then if you revert to an earlier 6.3 release, the storage system will attempt to write as many of the paths as possible back to the registry for the `cifs.home_dir` option. If some of the paths fail to be written to the registry, you might need to manually transfer the paths from the `cifs_homedir.cfg` file to the `cifs.home_dir` option. The `cifs_homedir.cfg` file remains intact during reversion.

Creating Diskettes

A

About this appendix This appendix describes how to create the following diskettes from a UNIX or Windows client:

- ◆ System boot diskettes
- ◆ A system firmware diskette

When to create diskettes

You do not need to create diskettes for F87, F825, F880, FAS200 series, FAS900 series, FAS3000 series, or NearStore series storage systems. These systems do not have a floppy drive and use a CompactFlash card for firmware, diagnostics, and system software downloads.

For all other storage systems (F810, F820, and F840 systems):

- ◆ You need to create a set of system boot diskettes if you do not have boot diskettes for the release that you are running.
- ◆ You need to create a firmware upgrade diskette whenever you need to upgrade firmware.

Diskette requirements

For each diskette that you want to create, you must have a write-enabled, 1.44-MB, PC-formatted diskette.

The number of diskettes you need for each kind of software is shown in the following table.

Software	Number of diskettes required
System boot	8
Firmware	1

Topics in this appendix

This appendix discusses the following topics:

- ◆ [“Creating diskettes from NOW”](#) on page 160
- ◆ [“Creating diskettes from the CD-ROM”](#) on page 164

Creating diskettes from NOW

About creating diskettes from NOW

You can create diskettes from either a UNIX client or a Windows client. You use different procedures to create boot diskettes and firmware diskettes.

Creating a system boot diskette set from a UNIX client

From a UNIX client, to create a system boot diskette set from the NOW site, complete the following steps.

Step	Action
1	Log in to NOW.
2	Go to the NOW software download page for this release.
3	Copy the UNIX boot diskette files to your client.
4	Copy each boot diskette file to a separate diskette using the following command: <pre>dd if=file of=drivedev bs=512</pre> <i>file</i> is the name of the file. <i>drivedev</i> is the device name of the diskette drive. Result: A message listing file counts similar to the following example appears, confirming that the diskette has been written. <pre>2880+0 records in 2880+0 records out</pre>
5	Remove, write-protect, and label each diskette after creating it. Include the release number, diskette number, and storage system model. Example: Release 7.1, Disk 1 of 8, NetApp F840 Note_____ It is important to record the storage system model on the diskette label, especially if you have different models on site. Systems in different model series use different boot files. _____

Creating a system boot diskette set from a Windows client

From a Windows client, to create a system boot diskette set from the NOW site, complete the following steps.

Step	Action	
1	If...	Then...
	You do not have RaWrite on your Windows client	Complete the procedure described in “ Downloading the RaWrite utility to a Windows client ” on page 162, then go to Step 2 of this procedure.
	You have RaWrite on your Windows client	Go to Step 2 of this procedure.
2	Log in to NOW.	
3	Go to the NOW software download page for this release.	
4	Copy the boot diskette files to your client.	
5	If you are running...	Then...
	The Windows version of RaWrite	Double-click RAWRITEWIN.EXE.
	The DOS version of RaWrite	Enter the following command at the Windows client command line: rawrite
6	For each diskette file in the system boot diskette set, follow the RaWrite prompts. As appropriate, enter the path name of one of the boot files that you downloaded and the drive letter for the diskette. Result: RaWrite displays confirmation messages and writes the diskette.	

Step	Action
7	Remove, write-protect, and label each diskette after creating it. Include the release number, diskette number, and storage system model. Example: Release 7.1, Disk 1 of 8, NetApp F840 Note _____ It is important to record the storage system model on the diskette label, especially if you have different models on site. Systems in different model series use different boot files. _____

Downloading the RaWrite utility to a Windows client

You need the RaWrite utility on your Windows client to extract the system boot diskette software and write it to the diskette. To download the RaWrite utility to your Windows client, complete the following steps.

Step	Action
1	Log in to NOW.
2	Go to the Tools and Utilities page at the following URL: http://now.netapp.com/NOW/cgi-bin/toolsutils/ .
3	Find either the DOS or Windows RaWrite utility.
4	Copy the RaWrite utility file to your client.
5	If the file ends in .zip, unzip the file by double-clicking it.

Creating a system firmware diskette from either a UNIX or Windows client

To create a system firmware diskette from the NOW site, complete the following steps.

Step	Action
1	Log in to NOW.
2	Go to the System Firmware Release Table at http://now.netapp.com/NOW/tools/releasetable.shtml .
3	Click system firmware–related links until you see a page that displays firmware versions.

Step	Action
4	Select system firmware 4.3 or later.
5	Follow the directions on the download page for your firmware.
6	Remove, write-protect, and label the diskette after creating it. Include the firmware revision, diskette number, and storage system model. Example: Firmware 4.3, Disk 1 of 1, NetApp F840 Note_____ It is important to record the storage system model on the diskette label, especially if you have different models on site. Systems in different model series use different boot files. _____

Creating diskettes from the CD-ROM

About creating diskettes from the CD-ROM

You can create diskettes from either a UNIX client or a Windows client. From a Windows client, you use different procedures to create boot diskettes and system firmware diskettes.

Creating boot or firmware diskettes from a UNIX client

From a UNIX client, to create boot or firmware diskettes from a file on the Data ONTAP CD-ROM, complete the following steps.

Step	Action
1	Mount the Data ONTAP CD-ROM on your client.
2	Use the /files.htm file to locate the desired firmware or system boot software files.
3	Copy each desired firmware or boot file to a diskette using the following command: <pre>dd if=file of=drivedev bs=512</pre> <i>file</i> is the name of the file. <i>drivedev</i> is the device name of the diskette drive. Result: A message listing file counts similar to the following appears, confirming that the diskette has been written. <pre>2880+0 records in 2880+0 records out</pre>
4	Remove, write-protect, and label each diskette after creating it. Include the firmware revision or Data ONTAP release number, diskette number, and storage system model. Example: Release 7.1, Disk 1 of 8, NetApp F840 Note_____ It is important to record the storage system model on the diskette label, especially if you have different models on site. Systems in different model series use different boot files. _____

Creating a boot diskette set from a Windows client

From a Windows client, to create a boot diskette set from files on the Data ONTAP CD-ROM, complete the following steps.

Step	Action
1	Insert the Data ONTAP CD-ROM into your CD-ROM drive.
2	Open the Utils folder.
3	Open the RaWrite folder.
4	Double-click Rawrite.exe.
5	For each diskette file in the system boot diskette set, follow the RaWrite prompts, supplying an image source and the drive letter for the formatted diskette that will be one of the system boot diskettes. For the image source, enter the following path name: X:\ontap\Release\Type\bootN <i>X</i> is your CD-ROM drive letter. <i>Release</i> is the release number on the front of your CD-ROM. <i>Type</i> is x86 for an F810, F820, or F840 storage system. <i>N</i> is the number of the diskette you are creating. Result: RaWrite displays confirmation messages and writes the diskette.
6	Remove, write-protect, and label each diskette after creating it. Include the release number, diskette number, and storage system model. Example: Release 7.1, Disk 1 of 8, NetApp F840 Note_____ It is important to record the storage system model on the diskette label, especially if you have different models on site. Systems in different model series use different boot files. _____

Creating a firmware diskette from a Windows client

From a Windows client, to create a firmware diskette from a file on the Data ONTAP CD-ROM, complete the following steps.

Step	Action
1	Insert the Data ONTAP CD-ROM into your CD-ROM drive.
2	Open the CD-ROM and use the files.htm file to locate the desired firmware.
3	Copy the desired file to the diskette.
4	Remove, write-protect, and label each diskette after creating it. Include the firmware revision, diskette number, and storage system model. Example: Firmware 4.3, Disk 1 of 1, NetApp F840 Note_____ It is important to record the storage system model on the diskette label, especially if you have different models on site. Systems in different model series use different boot files. _____

Why you should convert resource.frk files

When you upgrade a storage system to Data ONTAP 7.1 from a version of Data ONTAP prior to 6.0, any Macintosh client using Thursby Software Systems DAVE software starts to use the new Windows streams feature. This means that the information that was formerly stored in special resource.frk files, such as desktop icons for programs and files, is lost.

You can avoid losing this information by using a Windows conversion utility called `na_macstream` to make resource.frk files use the Windows streams feature.

When you use na_macstream

It is recommended to use `na_macstream` after you install Data ONTAP 7.1, but before Macintosh clients running DAVE software start to access their files.

If you did not run `na_macstream` before users accessed files, see “[Converting resource.frk files using na_macstream](#)” on page 169 for information on how to replace these files using the `/overwrite` option of the `na_macstream` command.

Note

This is a one-way conversion. After completing the conversion, you cannot go back to the original resource.frk file method if you want to revert to an earlier version of Data ONTAP.

Where na_macstream is located

The `na_macstream` conversion utility is in the `utils\na_macstream` directory on the Data ONTAP CD-ROM for this release.

The `na_macstream` utility is also on the NOW site on the Tools & Utilities page.

Copying na_macstream

To use na_macstream, you must copy na_macstream from the CD-ROM or the NOW site to a Windows host that is attached to the upgraded storage system.

Copying na_macstream from the CD-ROM: To copy na_macstream from the CD-ROM, complete the following steps.

Step	Action
1	Insert the CD-ROM into the Windows host's CD-ROM drive.
2	Copy the na_macstream utility from the utils\na_macstream directory. Note _____ Also included in the directory is a readme.txt file. Check this file for last-minute information about na_macstream. _____

Copying the utility from NOW: To copy the na_macstream utility from NOW, complete the following steps.

Step	Action
1	Log in to NOW.
2	Click Service & Support.
3	Click Download Software.
4	Click Tools & Utilities.
5	Locate the na_macstream utility and click the link to download the software to a directory on the Windows NT client.

Guidelines for using na_macstream

Follow these guidelines when using na_macstream:

- ◆ You must specify at least one path (file or directory) and only one action (/list or /upgrade).
- ◆ You can use the root of the path that maintains Macintosh users' files to convert all files.
If a directory is specified, all files and directories in it are converted.
- ◆ You can abbreviate all options.
- ◆ You can use options in any order.

**Converting
resource.frk files
using
na_macstream**

To convert resource.frk files to use streams, complete the following step.

Step	Action
1	From a Windows host connected to the storage system that has been upgraded to Data ONTAP 7.1, enter the following command from the MS-DOS Prompt window: <pre>na_macstream path [path...] {/list /upgrade} [/verbose] [/overwrite] [/stop-on-err]</pre> <i>path</i> is the Universal Naming Convention (UNC) path to the storage system and file or directory that is to be converted. /list lists the names of resource.frk files to be converted to streams. /upgrade upgrades or converts resource.frk files to use streams. /verbose displays the name of each file as it is processed. /overwrite allows any existing streams to be overwritten with the data in the resource.frk files. This option is useful if you did not run na_macstream before users accessed files. /stop-on-err stops the program when it encounters an error. If you do not use this option, na_macstream attempts to continue even if errors are encountered.

**Example of the
na_macstream
conversion**

The following command converts all resource.frk files in the share mshare on the storage system engmix and displays each file as it is processed:

```
na_macstream \\engmix\mshare /upgrade /verbose
```


Understanding Release Family Upgrades

C

About this appendix	This appendix explains why differences between releases and release families might pose issues during upgrades.
Why there might be issues	Each release family introduces new features, some of which pose compatibility issues when upgrading from previous releases. Upgrade procedures are created to resolve these issues. Some issues are resolved automatically in the Data ONTAP software, but a few issues require manual configuration. When you upgrade and there are one or more intermediate release families between your source and target release, the latest release usually includes any automatic upgrade software included in previous releases (unless otherwise specified). However, you might need to review and resolve upgrade issues associated with intermediate release families before upgrading to the new release.
Differentiating among types of upgrades	Data ONTAP releases are identified by a combination of numerals and letters in three positions, x.y.z. When you upgrade from one Data ONTAP release to another, the upgrade can be <i>within</i> a release family or <i>between</i> release families. A release family is the set of releases that have the same feature release name denoted by the first two digits of the release number; for example, 6.5, 6.5R1, 6.5.1, 6.5.1P1, 6.5.1R1, and 6.5.2 are all releases within the same release family. Data ONTAP release types are defined in the Data ONTAP Release Model, available on the NOW site at the following address: http://now.netapp.com/NOW/products/ontap_releasemodel/index.shtml.
Upgrading within a release family	An upgrade <i>within</i> a release family is one in which the release number x.y.z does not change in the x or y components, only in the z components, when comparing the original to the target release. The following are examples of upgrades within release families: ◆ 7.0 to 7.0.0.1 ◆ 7.0.0.1 to 7.0.1 ◆ 7.1 to 7.1P1 Major changes are not usually introduced within a release family, and updates to system and disk firmware are not usually required when you upgrade. Rather, these releases usually include bug fixes and minor feature enhancements.

For these reasons, upgrades within a release family are usually simpler and involve less service disruption than upgrades between release families.

Upgrading between release families

An upgrade *between* release families is one in which the release number *x.y.z* changes in the *x* or *y* components from the original to the target release. The following are examples of upgrades between release families:

- ◆ 6.5.3H1 to 7.1
- ◆ 6.5 to 7.1
- ◆ 5.3 to 7.1

A new release family usually includes major changes in infrastructure and subsystems. When you upgrade from one release family to another, one or more of the following might have been introduced on your platform:

- ◆ Fundamental infrastructure changes—for example, changes to WAFL or RAID operation
- ◆ Version number changes requiring a filesystem upgrade—for example, in RAID, WAFL, Non-Volatile LOG (NVLOG) or Java subsystems
- ◆ New system firmware or disk firmware requirements

Such feature changes and requirements are cumulative between succeeding release families. You do not have to upgrade sequentially to each new release family—in other words, you can skip release families—but you must comply with the requirements of any intermediate release and you should be aware of any new system behavior introduced in an intermediate release. For example, if you are upgrading from 6.4.1 to 7.1, you must satisfy the upgrade requirements of the 6.5 release family.

For these reasons, upgrades between release families sometimes take longer, involve more steps, and interrupt storage system services longer than upgrades within a release family.

Optimizing Service Availability During Upgrades

D

About this appendix This appendix explains the factors affecting service availability during upgrades and provides details about service and protocol behavior.

About service availability during upgrades

Service availability during Data ONTAP upgrades can be optimized through planning and configuration. In some cases, upgrades can be completely nondisruptive from the clients' perspective.

The following factors can affect the availability of storage system services during upgrades:

- ◆ Whether the system being upgraded is a single node or a cluster partner
Clustered systems are designed to provide optimal service availability.
For more information, see [“Upgrade host requirements”](#) on page 6.
- ◆ The types of services licensed and their susceptibility to timeout errors
See [“Service and protocol considerations”](#) on page 174.
- ◆ Upgrading between release families or within a release family
Upgrades between release families involve more steps and are potentially more disruptive than upgrades within a release family.
For more information, see [“Resolving upgrade issues”](#) on page 16.
- ◆ Whether a system firmware upgrade is required
System firmware upgrades require a system halt and reboot. This can disrupt services in single storage system and standard cluster upgrades but do not affect services in nondisruptive cluster upgrades.
For more information, see [“Upgrading the system firmware”](#) on page 60.
- ◆ Whether disk firmware upgrades are required and what type of RAID protection applies to those disks
For more information, see [“Upgrading disk firmware”](#) on page 67.
- ◆ The types of applications in use and their susceptibility to timeout errors
The availability of client applications during upgrades depends on features, protocols, and configuration. See your application documentation for more information.

Note

All hardware and software upgrades in any storage solution are potentially at least somewhat disruptive to storage system services. Make sure that you review upgrade options carefully to determine the best method of upgrading for maintaining optimal service availability.

**Service and
protocol
considerations**

During an upgrade, the storage system must be rebooted (by issuing the reboot command or by initiating a cluster takeover and giveback) to load the new software. In general, services based on stateless protocols—such as NFS, FCP, and iSCSI—are less susceptible to service interruptions during upgrades than session-oriented protocols—such as CIFS, FTP, NDMP, and HTTP. Furthermore, services based on stateless protocols usually remain available during nondisruptive upgrades of clustered systems.

Note

If you are planning to perform a nondisruptive upgrade on a system with FCP or iSCSI hosts, see [“About nondisruptive cluster upgrades”](#) on page 8.

Stateless protocols usually include a timeout procedure. For example, if a message is sent and receipt is not acknowledged within a timeout period, a transmission error is assumed to have occurred. In a storage system environment, if the client’s timeout period is greater than the disruption period on the storage system (for example, the amount of time a reboot or cluster giveback takes), the client does not perceive a disruption of storage system services.

In session-oriented protocols, there is no concept of timeout to protect the service from disruption. If session-oriented storage system services are disrupted, state information about any operation in progress is lost and the user must restart the operation.

For stateless protocols: The following list summarizes the behavior of storage system services and stateless protocols during upgrades. See also *The Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products*, on the NOW site at the following address:

http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

- ◆ NFS hard mounts

No adverse behavior on the clients. Clients might receive some messages similar to NFS server not responding, retrying until the storage system reboots.

- ◆ NFS soft mounts

You are advised not to use soft mounts for NFS traffic.

When the host is fabric-attached to a storage system that is not clustered, the recommended value for these parameters is 255 seconds.

- ◆ FCP, Direct Attached

No adverse behavior on the clients if the reboot takes less than the user-configurable node and link timeouts on the following hosts:

- ❖ For Windows 2000 hosts, the host parameters are called `NodeTimeout`, `HLinkTimeout`, and `LinkTimeout`.
- ❖ For Solaris hosts, the host parameter is called `nodev-tmo`, which is set in the `/kernel/drv/lpfc.conf` file.

The recommended value for these parameters for non-multipathed FCP configurations is 255 seconds. (Check the NOW site for the latest values.)

- ◆ FCP, Fabric Attached

No adverse behavior on the clients if the reboot takes less than the user-configurable node timeouts on the following hosts:

- ❖ For Windows 2000 hosts, the parameter is called `NodeTimeout`.
- ❖ For Solaris hosts, the parameter is called `nodev-tmo`.

When the host is fabric-attached to a storage system that is not multipathed, the recommended value for these parameters is 255 seconds.

- ◆ iSCSI

No adverse behavior on the clients unless otherwise noted—see “[iSCSI targets unavailable after upgrading to Data ONTAP 7.1](#)” on page 18. Clients might receive some messages similar to `Server not responding`, retrying until the storage system reboots.

For session-oriented protocols: The following list summarizes the behavior of storage system services and session-oriented protocols during upgrades.

- ◆ CIFS

Client sessions are terminated. Inform users to end their sessions before you upgrade. To do so, issue the `cifs terminate -t` command before the cluster takeover or before the reboot, or issue the `reboot -t` command.

- ◆ FTP, NDMP, and HTTP

State is lost and the client user must retry the operation.

- ◆ Backups and restores

State is lost and the client user must retry the operation.

Attention

Do not initiate a backup or restore operation during or immediately before an upgrade. Doing so might result in data loss.

- ◆ Applications (for example, Oracle or Exchange)

Effects depend on the applications. For timeout-based applications, you might be able to change the timeout setting to longer than the Data ONTAP reboot time to minimize adverse effects.

Index

Symbols

/etc/disk_fw directory, automatic disk firmware upgrades 67
/etc/exports file
 new access rules in Data ONTAP 6.5.x 39
 reverting 154
/etc/quotas file, new options in Data ONTAP 6.1 and later 48
/etc/resolv.conf file, creating of 75
/etc/software directory
 listing the contents of 90
 upgrading from 87

Numerics

1-Gb and 2-Gb mode, FC-AL adapters configured for 26

A

access control
 changes after upgrading 30
 reversion issues 146
access rules, new in Data ONTAP 6.5.x 39
administrator account, changes in Data ONTAP 6.1.1 and later 49
aggr command, using after upgrading 31
aggregates, migrating to after upgrading 31
AT-FCX modules, reversion issues 146
AutoSupport analysis, using to detect disk firmware requirements 25, 71
availability, service 173
 See also nondisruptive cluster upgrades

B

background disk firmware upgrade 68, 69
backups, behavior during upgrade 175
boot diskettes
 about 159
 creating from CD-ROM on a UNIX client 164
 creating from CD-ROM on a Windows client

 165
 creating from NOW on a UNIX client 160
 creating from NOW on a Windows client 161
 RaWrite utility 161
bugs
 105365 35
 113823 24
 129083 43
 98952 36
 information about 3
Bugs Online 3
BURT 99524, 102606, 104365 152

C

case-sensitivity, changes for user names 30
CHAP (Challenge Handshake Authentication Protocol), iSCSI upgrade issues 43
CIFS 175
 changes to administrator account in Data ONTAP 6.1.1 and later 49
 DNS requirements for Windows 2000 controllers 75
 domain account requirements for Windows NT controller 75
 nondisruptive cluster upgrade limitations 10
 upgrade recommendations 38
Cluster Configuration Checker 8
clustered storage systems
 about upgrading 8
 choosing upgrade procedure 99
 factors affecting service availability 173
 new disk firmware update procedure in Data ONTAP 6.1.1 49
 nondisruptive upgrades 8, 101, 109, 117
 reverting with software-based disk ownership 147
 standard upgrades 11
 upgrading system firmware 62
CompactFlash cards
 32MB card limits 35
 required upgrade for systems with 35

- service partition format 36
- Computer Management tool, sharing vFiler directories with 44
- CPU activity, increased in DATA ONTAP 6.5x 39
- create_reserved option, changes after upgrading 28

D

- D patch releases 3
- Data ONTAP
 - release families, about 171
 - supported versions 4
- Data ONTAP releases, known issues
 - 6.1 family 48
 - 6.2 and earlier 26
 - 6.2 family 47
 - 6.4 family 34, 43, 46, 167
 - 6.4.1 36
 - 6.4.2 and earlier 36
 - 6.5 family 34
 - 7.0 family 18, 24
- DAVE client, new streams feature in 6.4 46, 167
- directories, reverting large 154
- disk activity, increased in Data ONTAP 6.5x 39
- Disk Drive & Firmware Matrix 72
- disk firmware
 - about upgrading 67
 - background upgrade 68, 69
 - cluster upgrade requirements 10
 - command for upgrading 72
 - detecting upgrade requirements 71
 - manual upgrades 72
 - new update procedure for clustered storage systems in Data ONTAP 6.1.1 49
 - obtaining updates 72
 - RAID type considerations 10, 68
 - requirements 5
 - standard upgrade 68, 70
 - upgrade issues 24
- disk shelves, upgrade issues 24
- Disk Support Matrix & Firmware 5
- disk_fw_update command 72
- diskettes. *See* boot diskettes, system firmware diskettes
- DiskShelf14mk2 AT shelves, reversion issues 146

DNS

- configuring 75
- requirements for Windows 2000 name server 75
- domain account requirements 75
- double parity (DP), RAID. *See* RAID-DP
- DS14 disk shelves
 - reversion issues 147
 - upgrade issues 24
- DS14Mk2 disk shelves
 - reversion issues 147
 - upgrade issues 24

F

- F810 storage systems
 - diskette requirements 159
 - system firmware requirements 60
 - upgrade issues 26
 - upgrade procedures, clustered storage systems (standard) 132
 - upgrade procedures, single storage systems 97
- F820 storage systems
 - diskette requirements 159
 - system firmware requirements 60
 - upgrade issues 26
 - upgrade procedures, clustered storage systems (standard) 132
 - upgrade procedures, single storage systems 97
- F825 storage systems
 - reversion issues 147
 - system firmware requirements 60
 - upgrade issues 27, 35, 36, 37
 - upgrade procedures, clustered storage systems (nondisruptive) 101, 109, 117
 - upgrade procedures, clustered storage systems (standard) 128
 - upgrade procedures, single storage systems 95
- F840 storage systems
 - diskette requirements 159
 - system firmware requirements 60
 - upgrade issues 26
 - upgrade procedures, clustered storage systems (standard) 132
 - upgrade procedures, single storage systems 97

- F87 storage systems
 - diskette requirements 159
 - system firmware requirements 61
 - upgrade issues 37
 - upgrade procedures 97
- F880 storage systems
 - reversion issues 147
 - system firmware requirements 60
 - upgrade 27, 35, 36
 - upgrade issues 37
 - upgrade procedures, clustered storage systems (nondisruptive) 101, 109, 117
 - upgrade procedures, clustered storage systems (standard) 128
 - upgrade procedures, single storage systems 95
- FAS200 series storage systems
 - system firmware requirements 60
 - upgrade issues 24
 - upgrade procedures, clustered storage systems (nondisruptive) 101, 109, 117
 - upgrade procedures, clustered storage systems (standard) 124
 - upgrade procedures, single storage systems 92
- FAS3000 series storage systems
 - reversion issues 146
 - upgrade procedures, clustered storage systems (nondisruptive) 101, 109, 117
 - upgrade procedures, clustered storage systems (standard) 124
 - upgrade procedures, single storage systems 92
- FAS900 series storage systems
 - reversion issues 146, 147
 - system firmware requirements 60
 - upgrade issues 27, 35, 36, 37
 - upgrade procedures, clustered storage systems (nondisruptive) 101, 109, 117
 - upgrade procedures, clustered storage systems (standard) 128
 - upgrade procedures, single storage systems 95
- FC-AL adapters, upgrading storage systems with 26
- FCP protocol
 - behavior during upgrade 175
 - nondisruptive cluster upgrades 9
- Fibre Channel

- disk firmware. *See* disk firmware
- Fibre Channel-Arbitrated Loop. *See* FC-AL
- firmware, disk. *See* disk firmware
- firmware, system. *See* system firmware
- FlexCache volumes, reversion requirements 146
- FlexClone, license requirements 20
- flexible volumes. *See* FlexVol volumes
- FlexVol volumes
 - changing root volume after upgrading 28
 - migrating data to after upgrading 28
 - reversion requirements 146
 - SnapMirror requirements 13
 - vol status output changes 31
- FTP
 - default behavior change 44
 - nondisruptive cluster upgrade limitations 10
 - protocol behavior during upgrade 175
 - redirection default change 44

H

- hard links, moving before revert 154
- home directory paths, reverting 157
- hosts, upgrade 6
- HTTP server
 - copying upgrade files without installing 87
 - deleting downloaded files 91
 - installing upgrade files 88
 - listing downloaded software 90
 - requirements 7
 - upgrading from 6, 87
- HTTP, protocol behavior during upgrade 175

I

- inodes, reverting with larger 152
- intermediate upgrades 16, 26, 36
- iSCSI
 - nondisruptive cluster upgrades 9
 - protocol behavior during upgrade 175
 - upgrade issues with CHAP 43

L

- large files, transferring after upgrade 29
- license codes

- entering new 135
- FlexClone requirements 20
- SnapLock requirements 20
- local administrator account, changes in Data ONTAP 6.1.1 and later 49

M

- Macintosh client, new streams feature in 6.4 46, 167
- Microsoft Management Console. *See* MMC
- MMC, sharing vFiler directories with 44

N

- na_macstream utility 46, 167
- NDMP
 - nondisruptive cluster upgrade limitations 10
 - protocol behavior during upgrade 175
- NDU. *See* nondisruptive cluster upgrades
- NearStore series storage systems. *See* R100 storage systems, R150 storage systems, R200 storage systems
- NFS
 - access rules, new in Data ONTAP 6.5.x 39
 - export bug fix 43
 - protocol behavior during upgrade 174
- nfs.export.neg.timeout option 43
- nondisruptive cluster upgrades
 - about 8
 - limitations 10
 - requirements 99
 - system firmware upgrades 62
 - upgrade procedures 101, 109, 117
- nondisruptive disk firmware upgrade *See* background disk firmware upgrade.

O

- Open Systems SnapVault Agent, upgrade requirements 37
- optimizing service availability 174
 - See also* nondisruptive cluster upgrades

P

- P patch releases 3
- partition format, CompactFlash cards 36
- patch releases 3
- privileges
 - changes to user accounts after reverting 146
 - changes to user accounts after upgrading 30
- protocols, effects of upgrading on 174
- pseudo-root directories, vFiler 44

Q

- qtree replication with SnapMirror
 - transferring large files after upgrading 29
 - upgrade recommendations 12, 30
- quota initialization, avoid before upgrade 22, 38
- quotas
 - new options in Data ONTAP 6.1 and later 48
 - turned off during upgrade 32

R

- R100 NearStore systems
 - upgrade procedures 95
- R100 storage systems
 - system firmware requirements 60
 - upgrade issues 35, 36, 37
- R150 NearStore systems
 - upgrade procedures 95
- R150 storage systems
 - system firmware requirements 60
 - upgrade issues 35, 36, 37
- R200 NearStore systems
 - upgrade procedures 95
- R200 storage systems
 - system firmware requirements 60
 - upgrade issues 27
- raid.reconstruct_speed, new options in Data ONTAP 6.2 and later 47
- RAID4
 - disk firmware upgrade requirements 10, 68
 - group size limits 40
 - new options in Data ONTAP 6.2 and later 47
 - reversion errors 155
- RAID-DP

- disk firmware upgrade requirements 10, 68
 - upgrading from RAID4 41
- RaWrite utility
 - creating boot diskettes 161
 - obtaining 162
- redirection, FTP default behavior change 44
- release families
 - about 171
 - upgrade issues 16
 - See also* Data ONTAP release families, known issues
- requirements
 - Data ONTAP 4
 - disk firmware 5
 - diskette-based storage systems 159
 - domain accounts 75
 - Open Systems SnapVault Agent versions 37
 - SnapMirror upgrades 12
 - system firmware upgrades 60
 - upgrading 1, 74
- resource.frk file conversion, 6.4 upgrade 46, 167
- restores, behavior during upgrade 175
- restoring data from tape
 - upgrade recommendations 12, 30
- reverting
 - /etc/exports file 154
 - about 140
 - access control 146
 - contacting technical support 140
 - FlexCache volumes 146
 - FlexVol volumes 146
 - general guidelines 141
 - hard links, moving 154
 - home directory paths 157
 - large directories 154
 - larger inodes, with 152
 - pre-6.3 issues 157
 - RAID-related errors 155
 - root volumes 145
 - SecureAdmin 147
 - SnapLock volumes, removing 143, 156
 - SSL 147
 - subdirectories, deleting 154
 - upgrade must complete before 142
 - user account privileges 146

- vFiler units, stale file handles 157
 - vol status command output 146
- rolling upgrade. *See* nondisruptive cluster upgrades
- root volume
 - changing after upgrading 28
 - reversion requirements 145

S

- Secure Socket Layer. *See* SSL
- SecureAdmin
 - changes after upgrading 32
 - reverting 147
- Security ID (SID)-based quotas, new in Data ONTAP 6.1 and later 48
- serial number file, creating 136
- service partition format, CompactFlash cards 36
- services, effects of upgrading on 174
- sharing vFiler pseudo-root directories 44
- SnapDrive for Windows, upgrade issues 21
- SnapLock
 - removing volumes before reverting 143, 156
- SnapLock, license requirements 20
- SnapMirror
 - delete pre-6.1 snapshots 49
 - identifying source and destination volumes 13
 - upgrade requirements 12
 - volume type changes 13
- SnapMover, reversion issues 147
- snapshots, pre-6.1 should be deleted 49
- SnapVault
 - transferring large files after upgrading 29
 - upgrade recommendations 12, 30
- software command
 - about 87
 - delete subcommand 91
 - get subcommand 87
 - install subcommand 88
 - list subcommand 90
- software-based disk ownership, reverting 147
- space reservations settings, ignored after upgrade to Data ONTAP 6.4 and later 46
- space reservations, changes after upgrading 28
- SSL
 - must be reenabled after upgrading 32

- reverting 147
- standard cluster upgrades
 - about 11
 - requirements 99
 - upgrade procedures 128
- subdirectories, deleting before revert 154
- system firmware
 - 4.2.2, known issues 35
 - 4.2.3, known issues 35
 - command for updating 61
 - determining upgrade needs 61
- diskettes
 - about 159
 - creating from CD-ROM on a UNIX client 164
 - creating from CD-ROM on a Windows client 166
 - creating from NOW 162
- obtaining updates 60
- requirements 4
- upgrade issues 35
- upgrade requirements 60
- upgrading for clustered storage systems 62
- upgrading for single storage systems 62

T

- traditional volumes
 - changing root volumes after upgrading 28
 - preserved after upgrading 28
 - SnapMirror requirements 13
 - space reservations changes after upgrading 28
 - vol status output changes 31
- Troika card, upgrading clustered storage systems with 65, 66

U

- UNIX client
 - copying system files from the CD-ROM 80
 - downloading system files to 79
 - installing upgrade files 80
 - mounting the storage system on 78
 - requirements 7
 - upgrading from 6, 78
- UNIX IDs, used to track quotas in Data ONTAP 6.1

- and later 48
- update-flash command 61
- upgrade hosts 6
- Upgrade Utility 86
- upgrading
 - about 1
 - clustered storage systems 8
 - copying system files from UNIX host 80
 - copying system files from Windows PC 83, 84
 - disk firmware 24, 67
 - diskette requirements 159
 - downloading to clustered storage systems 99
 - downloading to single storage systems 95
 - FC-AL adapters, storage systems with 26
 - from a UNIX client 6, 78
 - from a Windows client 6, 82
 - from an HTTP server 6, 87
 - from client hosts 6
 - installing system files 77
 - intermediate upgrades 16, 26, 36
 - new license codes, entering 135
 - nondisruptive. *See* nondisruptive cluster upgrades
 - Open Systems SnapVault Agent requirements 37
 - protocol behavior during upgrade 174
 - qtree replication with SnapMirror 12
 - quotas turned off 32
 - release family issues 16, 171
 - requirements 4, 74
 - restoring data from tape 12
 - serial number file, creating 136
 - service availability during upgrade 173
 - SnapDrive for Windows requirements 21
 - SnapMirror 12
 - SnapVault 12
 - system firmware issues 35
 - system firmware requirements 60
 - volume type changes 28
- user accounts
 - privilege changes after reverting 146
 - privilege changes after upgrading 30
- user names, changes to case sensitivity 30
- user quotas, new options in Data ONTAP 6.1 and later 48

V

- vFiler migration, SnapMover reversion issues 147
- vFiler units
 - sharing pseudo-root directories 44
- vFiler units, stale file handles on reversion 157
- virus scanning, timeout option change 32
- vol commands, using after upgrading 31
- vol status command, output changes 31, 146
- volume types, changes in Data ONTAP 7.0 28
- volume, root. *See* root volumes
- volumes, flexible. *See* FlexVol volumes
- volumes, traditional. *See* traditional volumes
- vscan options, changes after upgrading 32

W

- Windows client
 - copying and installing system files from the CD-ROM 84
 - downloading and installing system files 83
 - mapping the storage system to a drive 82
 - requirements 7
 - upgrading from 6, 82
- Windows NT domain controller, requirements for CIFS users 75
- Write-Once-Read-Many (WORM). *See* SnapLock.

