

Data ONTAP® 8.0 7-Mode **Upgrade Guide**

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

Part number 210-04999_A0

Updated for Data ONTAP 8.0.1 on 18 November 2010

Contents

Copyright information	9
Trademark information	11
About this guide	13
Audience	13
Accessing Data ONTAP man pages	14
Terminology	14
Where to enter commands	16
Keyboard and formatting conventions	17
Special messages	18
How to send your comments	18
Planning your upgrade	19
Using the Upgrade Advisor to plan your upgrade	19
Upgrade process overview	21
Recommendations for all systems upgrading to this release	22
Upgrade host requirements	23
Requirements when upgrading from a Windows or UNIX client using the CIFS or NFS protocols	23
Requirements when upgrading from an HTTP server	24
Upgrade requirements for SnapMirror	24
Why you must plan for SnapMirror upgrades	25
SnapMirror synchronous and asynchronous mode during upgrade	25
Upgrade requirements for systems mirroring each other	26
Release family upgrade requirements	26
Different types of upgrades	26
Upgrades between release families	27
Upgrades within a release family	27
Required intermediate upgrades	28
Nondisruptive upgrade requirements	28
When to use nondisruptive high-availability upgrades	28
When not to use nondisruptive upgrades	29
Requirements for nondisruptive upgrades on all systems	30

Requirements for nondisruptive upgrades on systems with deduplicated volumes	31
Standard upgrade requirements	32
Evaluating upgrade issues	32
Issues to resolve before upgrading to the Data ONTAP 8.0 release family	33
Changes to behavior in the Data ONTAP 8.0 release family	36
Issues to resolve before upgrading from releases earlier than Data ONTAP 7.3	38
Behavior changes when upgrading from releases earlier than Data ONTAP 7.3	39
Preparing for the upgrade	43
Verifying system requirements	44
Ensuring that your system supports the target Data ONTAP release	45
Ensuring that there is adequate free space in every volume containing LUNs	45
Deduplication upgrade requirements	46
Determining the required firmware for your disks	46
Determining the required firmware for your disk shelves	46
Enabling DNS with Windows 2000 name server addresses	47
Verifying that you have a domain account	47
HA Configuration Checker	47
Preparing for nondisruptive upgrades	48
Preparing for nondisruptive upgrades on systems with VMware ESX server hosts	50
Determining system capacity and space guarantees before upgrading to Data ONTAP 7.3 or later	51
Using the aggrSpaceCheck tool to prepare your upgrade to Data ONTAP 7.3 or later	52
Reconfiguring IPv4 before upgrading	53
Obtaining Data ONTAP software images	55
Obtaining images for HTTP servers	56
Copying the software image to the HTTP server	56
Copying software images from the HTTP server without installing the images	56
Obtaining images for UNIX clients	57

Mounting the storage system on your client	58
Obtaining software images	58
Obtaining images for Windows clients	59
Mapping the storage system to a drive	59
Obtaining software images	60
Managing files in the /etc/software directory	61
Installing Data ONTAP software images	63
Installing software images from an HTTP server	64
Installing software images from the /etc/software directory	67
Downloading and rebooting new Data ONTAP software	71
Upgrading in a SnapMirror environment	71
Upgrading nondisruptively in a SnapMirror environment	72
Upgrading HA configurations from an earlier release family nondisruptively	73
Upgrading HA configurations within a release family nondisruptively	78
Upgrading HA configurations using the standard method	81
Upgrading single systems	84
Updating firmware	87
System firmware updates	87
Automatic BIOS system firmware updates	88
Updating system firmware nondisruptively	88
Updating system firmware using the standard method	91
Disk firmware updates	92
How disk firmware is updated	92
Service availability during disk firmware updates	93
Detecting outdated disk firmware	95
When to update disk firmware manually	96
Command for updating disk firmware	96
Disk shelf firmware updates	97
How disk shelf firmware is updated	97
Service availability during disk shelf firmware updates	98
Detecting outdated disk shelf firmware	99
Updating disk shelf firmware manually	100
Updating ACP firmware	102
Service Processor firmware updates	103
Using the Data ONTAP CLI to update the SP firmware	104
Using the SP CLI to update the SP firmware	104

RLM firmware updates	105
Requirements for RLM firmware version 4.0 and later	105
Using the Data ONTAP CLI to update the RLM firmware	106
Using the RLM CLI to update the RLM firmware	108
RLM firmware update problems	110
BMC firmware updates	111
Detecting outdated BMC firmware	112
Updating BMC firmware nondisruptively	113
Updating BMC firmware using the standard method	115
PAM II firmware updates	116
Reversion to a previous release	117
General guidelines for reverting from the Data ONTAP 8.0 release family	118
Guidelines for reverting systems with SnapMirror enabled	119
Order for SnapMirror system reversions	119
Preservation of SnapMirror relationships after reversion	119
Issues when reverting from Data ONTAP 8.0	120
Disabling compression for SnapMirror transfers after downgrading to Data ONTAP 8.0	121
Reinstatement of in-order frame delivery after reversion	121
Requirements for reverting a system with SSDs attached	121
Retention of modified security settings	121
Reversion issues for Brocade switches in fabric-attached MetroCluster	122
Changes to the interface group configuration in the /etc/rc file	122
Reverting with VLANs and an IP address configured on the base interface	122
Enabling TOE after reverting from Data ONTAP 8.0	123
Downgrade of deduplicated volumes with increased maximum size to Data ONTAP 8.0	123
Reversion of deduplicated volumes with increased maximum size	124
Reverting a SnapMirror destination system with volumes that use deduplication or clone operations	124
Reverting systems when a FlexClone file or FlexClone LUN operation is in progress	124
Reverting when Kerberos Multi Realm support is enabled	125
Issues when reverting to Data ONTAP 7.2	126
FlexCache reversion limitations	127

Deduplication reversion limitations	127
SnapMirror and SnapVault restart checkpoints deleted during reversion ...	128
SnapVault licenses might need to be removed before reverting	128
SnapVault restore processes must be complete before reverting	128
Large NFSv4 ACLs removed when reverting from Data ONTAP 7.3	128
FPolicy reversion issue with file names having long extensions	129
Optimal service availability during upgrades	131
How upgrades impact service availability	131
Service and protocol considerations	132
Considerations for stateless protocols	132
Considerations for session-oriented protocols	133
Index	135

Copyright information

Copyright © 1994–2010 NetApp, Inc. All rights reserved. Printed in the U.S.A.

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

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

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

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

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

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

Trademark information

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

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

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

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

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

NetApp, Inc. NetCache is certified RealSystem compatible.

About this guide

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

This document describes how to upgrade storage systems that run Data ONTAP to the latest release of firmware and software.

To determine whether your system is supported for the latest version of Data ONTAP, see the *Data ONTAP 7-Mode Release Notes* and system requirements.

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

Next topics

[Audience](#) on page 13

[Accessing Data ONTAP man pages](#) on page 14

[Terminology](#) on page 14

[Where to enter commands](#) on page 16

[Keyboard and formatting conventions](#) on page 17

[Special messages](#) on page 18

[How to send your comments](#) on page 18

Audience

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

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

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

Accessing Data ONTAP man pages

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

About this task

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

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

Step

1. View man pages in the following ways:

- Enter the following command at the console command line:


```
man command_or_file_name
```
- Click the manual pages button on the main Data ONTAP navigational page in the FilerView user interface.

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

Terminology

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

Storage terms

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

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

Cluster and high-availability terms

cluster	<ul style="list-style-type: none"> In Data ONTAP 8.x Cluster-Mode, a group of connected nodes (storage systems) that share a global namespace and that you can manage as a single virtual server or multiple virtual servers, providing performance, reliability, and scalability benefits. In the Data ONTAP 7.1 release family and earlier releases, a pair of storage systems (sometimes called <i>nodes</i>) configured to serve data for each other if one of the two systems stops functioning. In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an <i>active/active configuration</i>. For some storage array vendors, <i>cluster</i> refers to the hardware component on which host adapters and ports are located. Some storage array vendors refer to this component as a <i>controller</i>.
HA (high availability)	<ul style="list-style-type: none"> In Data ONTAP 8.x, the recovery capability provided by a pair of nodes (storage systems), called an <i>HA pair</i>, that are configured to serve data for each other if one of the two nodes stops functioning. In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an <i>active/active configuration</i>.
HA pair	<ul style="list-style-type: none"> In Data ONTAP 8.x, a pair of nodes (storage systems) configured to serve data for each other if one of the two nodes stops functioning.

- In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an *active/active configuration*.

Where to enter commands

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

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

Note: Data ONTAP commands shown in this document are for Data ONTAP 8.x 7-Mode and the Data ONTAP 7.x release families. However, some of these commands might also be available at the nodeshell prompt on systems running Data ONTAP 8.x Cluster-Mode. See the *Data ONTAP Cluster-Mode Administration Reference* for more information.

- You can enter commands either at the system console or from any client computer that can obtain access to the storage system using a Telnet or Secure Shell (SSH) session.
In examples that illustrate command execution, the command syntax and output shown might differ from what you enter or see displayed, depending on your version of the operating system.
- You can use the FilerView graphical user interface.
For information about accessing your system with FilerView, see the *Data ONTAP 7-Mode System Administration Guide*.
- You can enter Windows, ESX, HP-UX, AIX, Linux, and Solaris commands at the applicable client console.
In examples that illustrate command execution, the command syntax and output shown might differ from what you enter or see displayed, depending on your version of the operating system.
- You can use the client graphical user interface.
Your product documentation provides details about how to use the graphical user interface.
- You can enter commands either at the switch console or from any client that can obtain access to the switch using a Telnet session.
In examples that illustrate command execution, the command syntax and output shown might differ from what you enter or see displayed, depending on your version of the operating system.

Keyboard and formatting conventions

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

Keyboard conventions

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

Formatting conventions

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

Special messages

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

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

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

How to send your comments

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

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

Planning your upgrade

Because new features are introduced in each release of Data ONTAP, you must understand new features and upgrade requirements, and evaluate how they might impact your current configuration. You are more likely to encounter issues if you are upgrading from a release earlier than the immediately previous version of Data ONTAP.

Note: It is a best practice to use the Upgrade Advisor to plan your upgrade. If you are able to generate an upgrade plan with the Upgrade Advisor, it is not necessary to follow the instructions in this guide to identify planning, preparation, and upgrade procedures. Nonetheless, you might find useful detail and related information in this guide that complements your Upgrade Advisor plan.

If you are not able to use the Upgrade Advisor, you should create your own upgrade plan using guidelines provided in this guide.

Next topics

[Using the Upgrade Advisor to plan your upgrade](#) on page 19

[Upgrade process overview](#) on page 21

[Recommendations for all systems upgrading to this release](#) on page 22

[Upgrade host requirements](#) on page 23

[Upgrade requirements for SnapMirror](#) on page 24

[Release family upgrade requirements](#) on page 26

[Nondisruptive upgrade requirements](#) on page 28

[Standard upgrade requirements](#) on page 32

[Evaluating upgrade issues](#) on page 32

Using the Upgrade Advisor to plan your upgrade

You should use the Upgrade Advisor tool (if it is available in your environment) to ensure that you have met the requirements for upgrading to the current release and to generate an upgrade plan.

Before you begin

To use the Upgrade Advisor tool, your system must meet the following requirements:

- It must have a valid support contract.
- It must be enabled to send AutoSupport messages to NetApp.

Attention: If your system does not meet these requirements, you should consult the *Release Notes* and *Upgrade Guide* for this Data ONTAP release to prepare a detailed upgrade plan.

About this task

The Upgrade Advisor is an online tool, available on the NOW site, that simplifies the process of planning Data ONTAP upgrades. When you submit your system identification and target release to the Upgrade Advisor, the tool compares AutoSupport data about your system to known requirements and limitations of the target release. Upgrade Advisor then generates an upgrade plan (and optionally a back-out plan) with recommended preparation and execution procedures.

To generate an upgrade plan, you must have identifying information for your system (hostname, system ID, or serial number) and you must have selected a target upgrade release. You can also select other options, including the following:

- Create a plan for an HA pair, including nondisruptive upgrades.
- Create a back-out plan.
- Compare upgrade scenarios.

For more information about the Upgrade Advisor, see the Upgrade Advisor Help screens.

Steps

1. Locate and record the system hostname, system ID, or serial number of your system by entering the following command at the command line:

```
sysconfig
```

The system identification information is near the top of the display.

2. From a Web browser, log in to the My AutoSupport home page on NOW at the following URL:
`http://now.netapp.com/NOW/asuphome/`
3. Click the **Launch My AutoSupport** link.
4. Enter the hostname, system ID, or serial number of your system when prompted.
5. Select the system(s) you wish to upgrade from those listed.
6. Select the latest AutoSupport record from the ASUPs row.
7. Select the **Upgrade Advisor** tab.
8. Select the Data ONTAP release to which you want to upgrade from the **Target Versions** menu.
9. Select the upgrade method and the level of detail you would like included in your upgrade plan.
10. Click **Continue** to generate your upgrade plan.

After you finish

When you have created your upgrade plan, begin the upgrade process as recommended.

It is not necessary to follow further instructions in this *Upgrade Guide* after you have generated an upgrade plan using Upgrade Advisor. Nonetheless, you might wish to consult this guide for details and additional information.

Related information

[Upgrade Advisor -- now.netapp.com/NOW/asuphome/](http://now.netapp.com/NOW/asuphome/)

Upgrade process overview

Before beginning to upgrade Data ONTAP software, you should plan the upgrade and familiarize yourself with the required steps.

Attention: SnapLock technology is not supported in the Data ONTAP 8.0 release family. If you have SnapLock Compliance volumes, SnapLock Enterprise volumes, or copies of SnapLock volumes on your system, *do not* upgrade to any Data ONTAP 8.0.x release.

1. Plan your upgrade by familiarizing yourself with requirements and issues before you upgrade. Plan to do the following:
 - Review the Release Notes for your Data ONTAP upgrade target release.
 - Understand any requirements for upgrading to the target release from your existing software.
 - Attention:** You should use the Upgrade Advisor tool (if available in your environment) to assess your upgrade conditions and generate an upgrade plan.
 - Create a back-out plan, in the unlikely event that you need to revert to the Data ONTAP release that was running on your system before the upgrade. You should contact technical support if you need to revert to a previous release of Data ONTAP.
 - Note any potential changes to your system after the upgrade.
 - If you have storage systems in an HA pair, select the appropriate upgrade method.
 - If your storage system is in a SAN environment, verify that all components of your SAN configuration are compatible with the upgraded Data ONTAP release by consulting *NetApp Interoperability Matrix* on the NOW site.
 - If you run the SnapMirror software, identify storage systems with destination and source volumes.
 - If you are running MetroCluster systems, verify that all MetroCluster components are compatible with the target release.
2. If necessary, perform any required preparatory procedures before upgrading to the new Data ONTAP release. Required procedures might include the following:
 - Resolving upgrade issues, including performing an intermediate upgrade
 - Ensuring that you have a current Snapshot copy of the root volume of any system being upgraded
 - Updating disk firmware
 - Updating disk shelf firmware
 - Upgrading storage system firmware

3. Obtain the appropriate software image from the NOW site.
Copy the image to your storage system or to an HTTP server on your network.
4. Install the Data ONTAP software image on your storage system.
Extract the system files from the software image you copied to your system.
5. Download the new Data ONTAP system files to the boot device.
The upgrade process is completed when your HA pair or single system reboots with the new version of Data ONTAP.

Related concepts

[Planning your upgrade](#) on page 19

[Updating firmware](#) on page 87

[Lack of SnapLock support in Data ONTAP 8.0](#) on page 33

[Obtaining Data ONTAP software images](#) on page 55

[Installing Data ONTAP software images](#) on page 63

[Downloading and rebooting new Data ONTAP software](#) on page 71

[Reversion to a previous release](#) on page 117

Related tasks

[Preparing for the upgrade](#) on page 43

Related information

[Upgrade Advisor -- *now.netapp.com/NOW/asuphome/*](#)

[Download Software -- *now.netapp.com/NOW/cgi-bin/software*](#)

[NetApp Interoperability Matrix -- *now.netapp.com/NOW/products/interoperability/*](#)

Recommendations for all systems upgrading to this release

You should follow these simple guidelines to ensure your storage system upgrade goes smoothly.

- Review the "Important cautions" section of the *Release Notes* for this Data ONTAP release. It contains important information that could affect the behavior of your system during and after upgrading.
- Upgrade during non-peak hours.
- Avoid performing a quota initialization prior to upgrading.
If a quota initialization is in process prior to upgrading, wait for the initialization to finish.

Attention: SnapLock technology is not supported in the Data ONTAP 8.0 release family. If you have SnapLock Compliance volumes, SnapLock Enterprise volumes, or copies of SnapLock volumes on your system, *do not* upgrade to any Data ONTAP 8.0.x release.

The Data ONTAP 8.0 operating system is much larger than in previous release families. When you use the `download` or `software` command to activate Data ONTAP 8.0.x software images on

your storage system boot device, be aware that the `download` process takes significantly longer to finish than on earlier releases. For most systems upgrading to Data ONTAP 8.0.x releases, the download process finishes in 20 to 60 minutes. During this period, your system continues to serve data, but the system console is unavailable.

If you require system access during the `download` process, you can set the `telnet.distinct.enable` option, which allows you to open a Telnet or SSH-interactive session while the `download` command is running separately on the console. For more information about alternative methods of accessing the storage system, see the *Data ONTAP 7-Mode System Administration Guide*.

Upgrade host requirements

An *upgrade host* is the client system or server from which you upgrade Data ONTAP. You can upgrade Data ONTAP from a Windows or UNIX client, or from an HTTP server.

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

- The NOW site
- Portable storage media (such as a CD-R or USB drive) containing Data ONTAP software images
- An HTTP server containing Data ONTAP software images

Note: Beginning with Data ONTAP 8.0, CD-ROMs are no longer available for Data ONTAP upgrades.

You can install Data ONTAP system files after you prepare the upgrade host.

Next topics

[Requirements when upgrading from a Windows or UNIX client using the CIFS or NFS protocols](#) on page 23

[Requirements when upgrading from an HTTP server](#) on page 24

Related concepts

[Installing Data ONTAP software images](#) on page 63

Requirements when upgrading from a Windows or UNIX client using the CIFS or NFS protocols

If the CIFS or NFS protocols are licensed on your storage system, you can upgrade from a Windows or UNIX client using those protocols. You must be able to administer the storage system from the UNIX or Windows client. This client is usually the storage system's administration (admin) host.

Any UNIX or Windows admin host client with a network connection can be used to obtain Data ONTAP software images and copy them to a storage system.

For information about admin hosts, see the *Data ONTAP 7-Mode System Administration Guide*.

Note: The iSCSI protocol provides limited CIFS functionality that is not sufficient to allow you to upgrade using this method. With this limited functionality, you cannot create shares on the storage system for the software image.

Requirements when upgrading from an HTTP server

To upgrade from an 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.

Using an HTTP server is a good choice in these circumstances:

- The storage system does not have a CIFS or NFS license.
- You want to distribute Data ONTAP upgrade packages to multiple storage systems.
- You want to use installation scripts.

For information about the console, see the *Data ONTAP 7-Mode System Administration Guide*.

Related concepts

[Obtaining images for HTTP servers](#) on page 56

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. If you upgrade the source volumes first, SnapMirror volume replication is disabled. To reenable 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 than or the same as that on the destination system.

The requirement to upgrade SnapMirror destination volumes first applies to both asynchronous and synchronous SnapMirror for volume replication.

The requirement does not apply to SnapMirror for qtree replication, SnapVault, or data restoration for tape using the `restore` command. However, 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.

For more information about running SnapMirror on storage systems configured for network-attached storage (NAS), see the *Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide*.

Next topics

[Why you must plan for SnapMirror upgrades](#) on page 25

[SnapMirror synchronous and asynchronous mode during upgrade](#) on page 25

[Upgrade requirements for systems mirroring each other](#) on page 26

Related tasks

[Upgrading in a SnapMirror environment](#) on page 71

Why you must plan for SnapMirror upgrades

When you upgrade Data ONTAP on systems with SnapMirror relationships, the order in which you upgrade the systems is critical. If you do not upgrade in the correct order, SnapMirror transfers might not work correctly.

A SnapMirror transfer is possible only when the destination system can read a Snapshot copy of the source system. Therefore, the destination system must be upgraded first, because the upgraded destination system is able to read the Snapshot copies of the earlier release. If the source system is upgraded first, the destination system might not be able to read the source Snapshot copies, leading to failed SnapMirror transfers.

SnapMirror creates restart checkpoints during transfers, which allow an interrupted transfer to be restarted. These restart checkpoints are deleted during the following operations:

- Upgrade operation
- Revert operation
- System controller head swap operation

Once the restart checkpoints are deleted for an incomplete SnapMirror transfer, the transfer needs to be performed again from the start.

Note: Performing a SnapMirror transfer from the start is not the same as reinitializing the SnapMirror relationship. As long as there is a common Snapshot copy between the SnapMirror source and destination volumes, the destination can be updated with incremental transfers.

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

SnapMirror synchronous and asynchronous mode during upgrade

When you upgrade Data ONTAP on a destination storage system running on a synchronous mirror, SnapMirror goes into asynchronous mode.

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

Related tasks

[Upgrading in a SnapMirror environment](#) on page 71

Upgrade requirements for systems mirroring each other

To upgrade Data ONTAP on storage systems that are mirroring volumes to each other, you must disable the mirror, upgrade each system, and reenble the mirror.

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 system 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 reenble the mirror.

Release family upgrade requirements

Each Data ONTAP release family introduces new features. Most 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.

Next topics

[Different types of upgrades](#) on page 26

[Upgrades between release families](#) on page 27

[Upgrades within a release family](#) on page 27

[Required intermediate upgrades](#) on page 28

Different types of upgrades

Data ONTAP upgrades can be *within* a release family or *between* release families.

An upgrade *within* a release family is one in which the release number x.y.z does not change in the x or y components, but only in the z components of the release number. The following are examples of upgrades within release families:

- 8.0 to 8.0.1
- 7.3 to 7.3.1
- 7.2 to 7.2.5

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. For example, an upgrade from 7.3.3 to 8.0.1 is an upgrade between release families.

For more information about Data ONTAP release families and types of releases, see the Data ONTAP Release Model.

Related information

[Data ONTAP Release Model -- now.netapp.com/NOW/products/ontap_releasemodel/post70.shtml](http://now.netapp.com/NOW/products/ontap_releasemodel/post70.shtml)

Upgrades between release families

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 file system upgrade—for example, in RAID, WAFL, nonvolatile log (NVLOG), or Java subsystems
- New system firmware

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 7.2.1 to the current 8.0 release, you must satisfy the upgrade requirements of the 7.3 and 8.0 release families.

Note: Major nondisruptive upgrades (nondisruptive upgrades between release families) are supported only to a release in a succeeding release family. For example, you can upgrade directly from Data ONTAP 7.2.7 to 7.3.3 using the nondisruptive method, but not to 8.0.1. In such a case, you must upgrade nondisruptively through an intermediate release.

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.

Related concepts

[Requirements for nondisruptive upgrades on all systems](#) on page 30

[Required intermediate upgrades](#) on page 28

Upgrades within a release family

Upgrades within a release family are usually simpler and involve less service disruption than upgrades between release families.

This is because major changes are not usually introduced within a release family. Rather, these releases usually include bug fixes and minor feature enhancements.

Required intermediate upgrades

If you want to upgrade nondisruptively from a 7.2.x release to an 8.0.x release, you must perform an intermediate upgrade (also known as a multi-hop upgrade) to the latest 7.3.x release before upgrading to the target 8.0.x release.

In addition, if you are running a Data ONTAP 7.2 release earlier than 7.2.3, you must perform a minor NDU to the latest 7.2.x release before performing an intermediate major NDU to the latest 7.3.x release.

Attention: After performing an intermediate upgrade, you must wait at least 10 minutes before proceeding to the final upgrade (or to an additional intermediate upgrade) to ensure that all upgrade processes have finished.

There are no requirements for intermediate upgrades using the standard method (when you can schedule system downtime).

Nondisruptive upgrade requirements

Nondisruptive upgrades do not require downtime, and are available on some HA configurations.

In a nondisruptive upgrade (NDU), high-availability technology allows 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 HA node (storage system that is part of a high-availability relationship). Because the partner node fulfills service requests during the “failed” system’s upgrade, no disruption in service is experienced by the clients.

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

Next topics

[When to use nondisruptive high-availability upgrades](#) on page 28

[When not to use nondisruptive upgrades](#) on page 29

[Requirements for nondisruptive upgrades on all systems](#) on page 30

[Requirements for nondisruptive upgrades on systems with deduplicated volumes](#) on page 31

When to use nondisruptive high-availability upgrades

You can use the nondisruptive upgrade method on HA configurations that meet certain Data ONTAP requirements. Nondisruptive upgrades are most appropriate when high availability of storage system services is critical.

You can use the nondisruptive method when one or more of the following is being performed:

- Upgrades to the Data ONTAP 8.0 release family from an immediately preceding release family (for example, from 7.3.1 to 8.0)

Note: You can upgrade nondisruptively to the 8.0 release family from any release in the Data ONTAP 7.3 family.

If you need to upgrade from the 7.2 release family, you can upgrade nondisruptively from Data ONTAP 7.2.5 or later to the most recent 7.3 release, then upgrade nondisruptively to 8.0.

- Data ONTAP upgrades within a release family (for example, from 7.3 to 7.3.1)
- System firmware updates
- Certain hardware upgrades

Note: See the *Data ONTAP 7-Mode High-Availability Configuration Guide* for more information about changing system hardware nondisruptively.

When not to use nondisruptive upgrades

You cannot use the nondisruptive upgrade method in all circumstances.

Upgrades might be disruptive if any of the following conditions are true:

- 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 finish, and these services must be disabled to use nondisruptive upgrades.
- You need to update firmware for AT-FC-based or AT-FC2-based disk shelves.
Client services might encounter delays accessing data when disk shelf firmware is updated to AT-FC or AT-FC2 modules. To prevent data loss, all session-oriented services must be terminated before you begin an update procedure.
- You need to update disk firmware and you have RAID4 aggregates on your system.
Standard disk firmware updates automatically take disks in RAID4 aggregates offline until the update is complete. Services and data are unavailable until they are back online.

Note: If you upgrade RAID protection to RAID-DP, disk firmware updates take place in the background and are nondisruptive.

For these conditions, standard upgrades are recommended.

Related concepts

[Disk shelf firmware updates](#) on page 97

[Disk firmware updates](#) on page 92

[Service availability during disk firmware updates](#) on page 93

Requirements for nondisruptive upgrades on all systems

You must ensure that your systems meet configuration and utilization requirements before beginning a nondisruptive upgrade process.

Attention: Be sure to use the Upgrade Advisor tool (if it is available in your environment) to help you determine nondisruptive upgrade requirements.

Major nondisruptive upgrades (nondisruptive upgrades between release families) to Data ONTAP 8.0 releases are supported from all Data ONTAP 7.3 releases.

Note: If you are running a release in the Data ONTAP 7.2 release family and you want to upgrade nondisruptively to Data ONTAP 8.0 or later, you must first upgrade to the latest Data ONTAP 7.3.x release.

Minor nondisruptive upgrades (nondisruptive upgrades within release families) are supported from all previous Data ONTAP 8.0 releases.

To use the nondisruptive upgrade procedure, your systems must meet the following configuration requirements:

- You must have an HA pair in which a partner controller takes over I/O during the upgrade process.
- Because failed disk drives prevent giveback operations and can introduce loop instability throughout the storage system, you must remove or replace all failed disk drives *before* beginning the nondisruptive upgrade.
- There should be no old core files in the `/etc/crash` directory.
- Your systems must be running the latest disk and disk shelf firmware *before* beginning the nondisruptive upgrade.
- If your system serves NFS clients, you must use hard mounts.

Attention: You should not use soft mounts when there is a possibility of frequent NFS timeouts, which can lead to disruptions during the upgrade process and possible data corruption.

- You must be able to open a terminal session to the console port of both controllers in an HA pair using one of the following methods:
 - Direct serial connection
 - A console server
 - The systems' remote LAN modules (RLMs), if available
 - The systems' Baseboard Management Controllers (BMCs), if available

Because network connections to the controllers are lost during the takeover and giveback operations performed during the nondisruptive upgrade, Telnet, SSH, or FilerView sessions will not work.

You should avoid exceeding maximum values for the following system elements on all platforms:

Element	Value (per storage controller)
FlexVol volumes	500 Note: The limit for FAS2040 systems is 200 FlexVol volumes. Up to 300 of the maximum number of FlexVol volumes for your platform can be enabled for deduplication.
Snapshot copies	No more than 10 times the number of FlexVol volumes
CPU utilization	No greater than 50%
Disk utilization	No greater than 50%

Related concepts

[Requirements for nondisruptive upgrades on systems with deduplicated volumes](#) on page 31

[Optimal service availability during upgrades](#) on page 131

[Considerations for stateless protocols](#) on page 132

[Required intermediate upgrades](#) on page 28

Related tasks

[Using the Upgrade Advisor to plan your upgrade](#) on page 19

Requirements for nondisruptive upgrades on systems with deduplicated volumes

You can perform major and minor nondisruptive upgrades when deduplication is enabled, provided that no more than 300 FlexVol volumes have deduplication enabled and that no deduplication operations are running during the Data ONTAP upgrade.

The total number of deduplicated and non-deduplicated FlexVol volumes must not exceed the total number of FlexVol volumes supported for nondisruptive upgrades on your system.

Nondisruptive upgrades cannot take place when deduplication operations are active. To ensure that no deduplication operations are active, you must take both of the following actions:

- If any deduplication operations are active, you must halt them until the Data ONTAP upgrade has completed.
- You must perform the Data ONTAP upgrade during a time period when deduplication operations are not scheduled to run.

You can use the `sis status` command to determine if the status of a deduplication is `Active` or `Idle`. On a system with deduplication enabled, the output of the `sis status` command is similar to the following:

```
Path           State    Status    Progress
/vol/v457     Enabled  Idle      Idle for 00:12:30
```

/vol/v458	Enabled	Idle	Idle for 00:12:30
/vol/v459	Enabled	Idle	Idle for 00:12:30
/vol/v460	Enabled	Idle	Idle for 00:12:30
/vol/v461	Enabled	Active	521 MB Scanned
/vol/v462	Enabled	Active	489 MB Scanned
/vol/v463	Enabled	Active	387 MB Scanned
/vol/v464	Enabled	Idle	Idle for 00:12:30

You can use the `sis stop` command to abort the active SIS operation on the volume and the `sis start` command to restart it.

For information about deduplication, see the *Data ONTAP 7-Mode Storage Management Guide* and the `sis(1)` man page.

Standard upgrade requirements

A standard upgrade can be performed on any HA pair, but downtime is required.

In a standard upgrade, downtime is required because the HA configuration is disabled and each node is updated. When the HA configuration is disabled, each node behaves as a single-node storage system; in other words, system services associated with the node are interrupted for as long as it takes the system to reboot.

You can also complete other maintenance tasks, such as system firmware and hardware, as part of the standard upgrade. These can also take place when the HA pair is disabled.

Evaluating upgrade issues

Every Data ONTAP release family has unique upgrade requirements that you must understand and resolve before you decide to upgrade. Depending on your version of Data ONTAP, you might have to upgrade to an intermediate release before upgrading to the current release.

Before you decide to upgrade, you need to understand the following:

- 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 7.2 family to the current 8.0 release, you must review and resolve upgrade issues associated with the 7.3 and 8.0 release families before upgrading to Data ONTAP 8.0 or later.

Next topics

[Issues to resolve before upgrading to the Data ONTAP 8.0 release family](#) on page 33

[Changes to behavior in the Data ONTAP 8.0 release family](#) on page 36

Issues to resolve before upgrading from releases earlier than Data ONTAP 7.3 on page 38

Behavior changes when upgrading from releases earlier than Data ONTAP 7.3 on page 39

Issues to resolve before upgrading to the Data ONTAP 8.0 release family

You must understand and resolve these issues before you upgrade to Data ONTAP 8.0 and later releases.

Next topics

Lack of SnapLock support in Data ONTAP 8.0 on page 33

IPv6 not supported in Data ONTAP 8.0 on page 35

MetroCluster upgrade requirements for Brocade switches on page 35

FlexCache origin volumes running Data ONTAP 10.0.3 are not supported on page 35

IPsec is not supported in Data ONTAP 8.0 on page 35

cfmode support change in Data ONTAP 8.0 on page 36

Lack of SnapLock support in Data ONTAP 8.0

Data ONTAP 8.0 release family does not support the SnapLock feature. Therefore, while upgrading the storage system to Data ONTAP 8.0 release family, special considerations must be given to the storage system running the SnapLock feature.

Note: For continued access to the SnapLock volumes, do not upgrade storage system with SnapLock volumes to Data ONTAP 8.0 release family.

If the storage system halts because you have SnapLock volumes on the system and you attempted to upgrade to Data ONTAP 8.0, contact technical support immediately.

Upgrade of storage system with SnapLock license to Data ONTAP 8.0

You cannot add new SnapLock licenses in a storage system running Data ONTAP 8.0 or later in the 8.0 release family. Previously installed SnapLock licenses are retained, but are disabled. You can view the SnapLock license using the `license` command; however, all operations that require a SnapLock license fail. Therefore, you cannot create new SnapLock volumes or aggregates in Data ONTAP 8.0 release family.

A storage system might have the SnapLock license enabled, without any SnapLock volumes or aggregates. In such a case, if you upgrade the storage system to Data ONTAP 8.0 release family, the SnapLock license is disabled automatically and an error message `snaplock.unsupported.version` is raised while the storage system is booting up.

The SnapLock license is enabled automatically when you boot with a Data ONTAP release that supports the SnapLock feature.

Upgrade of storage systems with SnapLock volumes and aggregates to Data ONTAP 8.0

In some versions of Data ONTAP 7.2.x and Data ONTAP 7.3.x release families with SnapLock volumes, you might be able to upgrade to Data ONTAP 8.0.1. However, because Data ONTAP 8.0.1 does not support SnapLock feature, your storage system will halt.

The storage system console displays a list of disks that contain the SnapLock volumes and aggregates. Following is an example of an error message on the storage system console:

```
Found SnapLock disk:v5.29 Use fcadmin device_map for shelf and
slot info This release does not support SnapLock.Remove the
SnapLock disks from this system.
Found SnapLock disk:v5.32 Use fcadmin device_map for shelf and slot info
This release does not support SnapLock.Remove the SnapLock disks from this
system This
release does not support SnapLock.
Halting the system!!!
To recover-boot with a release that supports SnapLock or unplug the
SnapLock disks.
```

If you reboot a storage system running Data ONTAP 8.0 release family with SnapLock disks and non-SnapLock disks (from an earlier release), the storage system will halt.

Connection of disks that contains SnapLock aggregates to a storage system with Data ONTAP 8.0

If you upgrade the storage system to Data ONTAP 8.0 release family with disks that contains SnapLock aggregates, the storage system halts in the early boot process.

If you connect disks that contains SnapLock aggregates to a storage system with Data ONTAP 8.0 release family, the storage system remains online. However, the storage system lists these disks in the broken disk pool and displays the `snaplock.disk.on.unsupported.version` error message on the storage system console. The SnapLock disks in the broken disk pool are protected from data corruption.

Note: The SnapLock disks cannot be used in any other aggregate. When the SnapLock disks are moved into the broken disk pool in Data ONTAP 8.0 release family, the ComplianceClock associated with these volumes will not get updated. This might result in ComplianceClock skew when you reattach the SnapLock disks to a storage system running a Data ONTAP release that supports the SnapLock feature.

Reboot of the storage system running Data ONTAP 8.0 release family after connecting disks with SnapLock aggregates

If you reboot the storage system running Data ONTAP 8.0 and later releases of Data ONTAP 8.0 after connecting disks that contain SnapLock aggregates, the reboot is successful. However, you might not be able to use these disks. The storage system lists these disks in the broken disk pool and displays the `snaplock.disk.on.unsupported.version` error message on the storage system console.

IPv6 not supported in Data ONTAP 8.0

If IPv6 is enabled on your storage system, be aware that it will be disabled automatically during the upgrade to the Data ONTAP 8.0 release family. If you want to upgrade to the Data ONTAP 8.0 release family, you should take additional steps to ensure IPv4 network connectivity *before* upgrading.

Attention: If you have configured IPv6 on your storage system in an environment that requires IPv6 networking, do not upgrade to the Data ONTAP 8.0 release family.

If a network interface is configured with both IPv4 and IPv6 addresses, the IPv6 addresses are ignored after upgrading to Data ONTAP 8.0, and network traffic is sent over IPv4 addresses only.

IPv6 values for the following configurations are ignored after the upgrade:

- IPv6 entries in the Network Status Monitor (NSM) are skipped while sending NSM notifications to clients. Therefore, NSM notifications are sent only to IPv4 clients.
- IPv6 addresses present in the `/etc/usermap.cfg` and `/etc/exports` files are ignored.
- The exports rules for loading the `/etc/exports` file skip the IPv6 addresses present in each export rule. The entire exports rule is not skipped; only the IPv6 addresses in the exports rule are skipped.

However, if you have configured any networking services over IPv6, you must reconfigure IPv4 for those services before upgrading.

MetroCluster upgrade requirements for Brocade switches

If your storage system is running Data ONTAP 8.0.1 and you want to upgrade the switches to Brocade Fabric OS 6.3.1c or later, you must first create single-loop storage zones.

Note: If the switches are running Fabric OS 6.1.1a, you must first upgrade them to Fabric OS 6.2.x before upgrading to Fabric OS 6.3.1c or later.

For more information about creating single-loop storage zones, see the *Brocade Switch Configuration Guide for Fabric-attached MetroClusters* available on the NOW site.

FlexCache origin volumes running Data ONTAP 10.0.3 are not supported

If you have FlexCache volumes backed by origin volumes on systems running Data ONTAP 10.0.3, you must upgrade the origin system to Data ONTAP 10.0.4 or later before upgrading the caching system to the Data ONTAP 8.0 release family.

If you attempt to connect or create a FlexCache volume backed by a volume on a system running Data ONTAP 10.0.3, the caching system will panic.

IPsec is not supported in Data ONTAP 8.0

If IPsec is enabled on your storage system, be aware that it will be disabled automatically during the upgrade to the Data ONTAP 8.0 release family.

cfmode support change in Data ONTAP 8.0

For Fibre Channel SAN configurations in Data ONTAP 8.0 and later releases, only `single_image` cfmode (cluster failover mode) is supported. If you are upgrading high-availability FC SAN systems from an earlier release and they are configured for any other cfmode, you must migrate them to `single_image` mode *before* upgrading to Data ONTAP 8.0 or later.

For detailed instructions about migrating to `single_image` mode, see *Changing the Cluster cfmode Setting in Fibre Channel SAN Configuration* on the NOW site.

Related information

Changing the cluster cfmode setting in Fibre Channel SAN configurations - http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/

Changes to behavior in the Data ONTAP 8.0 release family

You should be aware of these changes in Data ONTAP behavior that might occur if you upgrade to Data ONTAP 8.0 or later.

Next topics

[New time protocol requirements](#) on page 36

[Obsolete timed options visible in Data ONTAP 8.0 7-Mode](#) on page 37

[Special system files](#) on page 37

[New minimum root volume sizes](#) on page 37

[RLM over IPv6 not supported in the Data ONTAP 8.0 release family](#) on page 37

New time protocol requirements

Starting with Data ONTAP 8.0 7-mode, the Network Time Protocol (NTP) protocol is the only supported protocol for time synchronization. The `rtc` and the `rdate` protocols of the `timed.proto` option are obsolete and no longer take effect after you upgrade to Data ONTAP 8.0 7-mode or later.

If your system does not use NTP as the time-synchronization protocol, it will not keep accurate time after you upgrade it to Data ONTAP 8.0 7-mode or later. Problems can occur when the storage system clock is inaccurate.

If your system does not already use NTP as the protocol for time synchronization, immediately after upgrading to Data ONTAP 8.0 7-mode or later you must set `timed.proto` to `ntp`, set `timed.servers` to use time servers that support NTP, and ensure that `timed.enable` is set to `on`.

Note: After you set `timed.proto` to `ntp`, the setting remains in effect even if you revert back to a release prior to Data ONTAP 8.0 7-mode.

For information about how to synchronize the system time, see the *Data ONTAP 7-Mode System Administration Guide*.

Obsolete timed options visible in Data ONTAP 8.0 7-Mode

Starting with Data ONTAP 8.0, several `timed` options are obsolete although they remain visible in the CLI and can be modified.

The following `timed` options have no effect in Data ONTAP 8.0 7-Mode or later:

- The `timed.max_skew` option
- The `timed.sched` option
- The `timed.window` option
- The `rtc` and the `rdate` protocols of the `timed.proto` option

If you attempt to set the above options when the system is running Data ONTAP 8.0 7-Mode or later, they will have no effect. However, these settings will take effect after the system is reverted back to a release that supports the options.

Special system files

For storage systems upgraded from a release earlier than Data ONTAP 8.0, some system files exist in every volume of the system. You must not remove or modify these files unless technical support directs you to do so. These files enable you to restore LUNs in Snapshot copies if you revert to a release earlier than Data ONTAP 8.0.

The following system files are in the root level of every volume, including the root volume:

- `.vtoc_internal`
- `.bplusvtoc_internal`

New minimum root volume sizes

The minimum required size for root volumes has been increased for every system running Data ONTAP 8.0. The new minimum sizes are not enforced when you upgrade from an earlier release, but if you modify the root volume, it must conform to the new requirements. If your root volume does not meet the new requirements, you should increase its size as soon as you complete the upgrade procedure.

The root volume must have enough space to contain system files, log files, and core files. If a system problem occurs, these files are needed to provide technical support. For more information about root volumes and the new size requirements for your platform, see the *Data ONTAP 7-Mode System Administration Guide*.

RLM over IPv6 not supported in the Data ONTAP 8.0 release family

If you upgrade from the Data ONTAP 7.3 release family and your RLM is configured with IPv6, existing IPv6 configuration is automatically removed during the upgrade because IPv6 is not supported in the Data ONTAP 8.0 release family.

If your RLM configuration does not include static IPv4 or DHCP IPv4, you need to reconfigure the RLM by using the `rlm setup` command after upgrading to the Data ONTAP 8.0 release family.

For information about the RLM, see the *Data ONTAP 7-Mode System Administration Guide*.

Issues to resolve before upgrading from releases earlier than Data ONTAP 7.3

You must understand and resolve certain issues before you upgrade from releases earlier than Data ONTAP 7.3.

Next topics

[More free space required in Data ONTAP 7.3](#) on page 38

[License changes for the FlexCache feature](#) on page 38

[Disks offline in Windows 2008 after a standard upgrade](#) on page 38

More free space required in Data ONTAP 7.3

Data ONTAP 7.3 includes an improvement to free space accounting. As a result, existing FlexVol volumes reserve additional space, resulting in a loss of 0.5 percent of free space. Upgrading to Data ONTAP 7.3 or later from an earlier release causes existing FlexVol volumes to require more free space from their containing aggregates. If there is insufficient free space in an aggregate to satisfy the increased requirement from its FlexVol volumes, the space guarantee for one or more volumes in that aggregate might be disabled.

Related tasks

[Determining system capacity and space guarantees before upgrading to Data ONTAP 7.3 or later](#) on page 51

License changes for the FlexCache feature

If you are currently using the FlexCache feature, you need to take action to continue to use this feature when you upgrade to Data ONTAP 7.3 and later.

The current FlexCache license, `flex_cache`, has been replaced by a new license, `flexcache_nfs`. The old license is supported for the Data ONTAP 7.2 release family, but is not supported for Data ONTAP 7.3 and later. See your sales representative to install the new `flexcache_nfs` license if it is not already present on your system.

Attention: If you upgrade to Data ONTAP 7.3 or later and the new license is not installed, you will not be able to access data in FlexCache volumes after the upgrade. As soon as you install the new license, the FlexCache data will become accessible.

Disks offline in Windows 2008 after a standard upgrade

During a standard upgrade to Data ONTAP 7.3.3 and later releases, LUNs are assigned new revision numbers. Windows Server 2008 software interprets the LUNs with new revision numbers as new

disks and sets them offline; this status is shown in Windows 2008 management interfaces after the upgrade. Windows Server 2003 ignores the LUN revision number.

You can work around this problem using the nondisruptive upgrade method, which allows the LUNs to maintain their revision numbers. You can also bring the disks online after the upgrade using Windows disk management tools or SnapDrive functionality.

For more information, see the knowledgebase article *Disks show as offline in Windows 2008 after Data ONTAP upgrade* on the NOW site.

Related information

[Disks show as offline in Windows 2008 after Data ONTAP upgrade: now.netapp.com/Knowledgebase/solutionarea.asp?id=kb54672](http://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb54672)

Behavior changes when upgrading from releases earlier than Data ONTAP 7.3

You should be aware of several changes in Data ONTAP behavior that might occur if you upgrade from releases earlier than Data ONTAP 7.3.

Next topics

[The NetBackup application can no longer manage SnapVault relationships with NetApp data](#) on page 39

[Physical reallocation of volumes slows the reversion process](#) on page 39

[SnapMirror and SnapVault restart checkpoints deleted during upgrade](#) on page 40

[Deduplication requires additional free space in aggregates after upgrading](#) on page 40

[FPolicy compatibility issue in NFSv4 environments](#) on page 40

[Kerberos Multi Realm support](#) on page 41

The NetBackup application can no longer manage SnapVault relationships with NetApp data

Beginning with Data ONTAP 7.3, the use of Symantec NetBackup for configuring and managing SnapVault transfers between NetApp primary and secondary storage systems is no longer supported.

If you are currently using the NetApp SnapVault Management option from Symantec, you can migrate to NetApp Operations Manager or Protection Manager, or to management using the command-line interface (CLI). This option is not supported with Data ONTAP 7.3 and later releases. You can continue to use this option with Data ONTAP 7.2.x and earlier. You should check with Symantec about support for this option for NetBackup versions later than 6.5.

Physical reallocation of volumes slows the reversion process

Data ONTAP 7.3 and later releases support physical reallocation, which allows you to optimize the physical layout of volumes in an aggregate, leaving the virtual location of the volumes untouched.

However, once volumes have been physically reallocated, reverting to an earlier release family will take significantly longer.

For more information about physical reallocation, see the *Data ONTAP 7-Mode System Administration Guide*.

SnapMirror and SnapVault restart checkpoints deleted during upgrade

Starting with Data ONTAP 7.3, when you upgrade to Data ONTAP 7.3 or later, all aborted qtree SnapMirror and SnapVault transfers with restart checkpoints will restart from the beginning because all restart checkpoints will be deleted during the upgrade process.

Deduplication requires additional free space in aggregates after upgrading

If you use deduplication, you must ensure that there is adequate free space in the aggregates containing deduplicated volumes after upgrading to Data ONTAP 7.3 or later.

In earlier Data ONTAP releases, the deduplication fingerprint database was stored in the deduplicated volume. In Data ONTAP 7.3 and later releases, the deduplication fingerprint database is automatically moved to the containing aggregate when deduplication is run for the first time on a volume after an upgrade. Before running deduplication for the first time, you should ensure that the aggregate has free space that is at least 4 percent of the total data usage for all volumes in the aggregate that have deduplication enabled, in addition to 2 percent free space for FlexVol volumes. This enables additional storage savings by deduplicating any new blocks with those that existed before the upgrade.

If there is not sufficient space available in the aggregate, the deduplication operation fails with an error message.

During a deduplication failure, there is no loss of data and the volume is still available for read/write operations. However, depending upon the space availability in the aggregate, fingerprints of the newly added data might be lost.

If you receive a deduplication failure message, you should add space to the aggregate (depending on the limits of your configuration) and run deduplication again.

For example, if an aggregate contains 3 FlexVol volumes and each volume has 5 TB of data (1 TB is physical usage and 4 TB is deduplication savings), the total data in the aggregate amounts to 15 TB. In such a case, after the upgrade, 600 GB (4 percent of 15 TB) and 300 GB (2 percent of 15 TB) must be available in the aggregate and volumes respectively.

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

FPolicy compatibility issue in NFSv4 environments

If you are running an application that uses the FPolicy engine and the application is running in an NFSv4 environment, you should upgrade the application to support NFSv4.

Beginning in Data ONTAP 7.3, FPolicy supports NFSv4. Previously, FPolicy did not support NFSv4 and NFSv4 requests were not passed on to any FPolicy-based application.

Although FPolicy now supports NFSv4, the FPolicy-based application might not support NFSv4. If an application that does not support NFSv4 receives notice of NFSv4 file operations (such as file OPEN and CLOSE events), these file operations might appear as UNKNOWN events to the application and generate error messages.

To avoid these compatibility problems, you should upgrade any FPolicy-based applications to support NFSv4.

Kerberos Multi Realm support

If you upgrade to Data ONTAP 7.3.1 or later from an earlier release, Data ONTAP continues to use the old keytab file for UNIX-based KDCs (`/etc/krb5.keytab`). You should only use the new keytab file for UNIX-based KDCs (`/etc/UNIX_krb5.keytab`) if you reconfigure Kerberos after such an upgrade or configure Kerberos for the first time.

In Data ONTAP 7.3.1 and later releases, you can configure Data ONTAP to use both Active Directory and UNIX-based KDC types simultaneously. This configuration is sometimes referred to as a "Kerberos Multi Realm" configuration.

To support Multi Realm configurations, Data ONTAP uses two sets of principal and keytab files. For Active Directory-based KDCs, the principal and keytab files are `/etc/krb5auto.conf` and `/etc/krb5.keytab`, respectively, just as in releases prior to Data ONTAP 7.3.1. For UNIX-based KDCs, however, the principal and keytab files are `/etc/krb5.conf` and `/etc/UNIX_krb5.keytab`, respectively. So, starting with Data ONTAP 7.3.1, the keytab file for UNIX-based KDCs has changed from `/etc/krb5.keytab` to `/etc/UNIX_krb5.keytab`.

This change does not affect upgrades, however, because Data ONTAP continues to use the old keytab file (`/etc/krb5.keytab`) for UNIX-based KDCs if you upgrade from a release prior to Data ONTAP 7.3.1. You need only use the new keytab file for UNIX-based KDCs (`/etc/UNIX_krb5.keytab`) if you reconfigure Kerberos after such an upgrade or you configure Kerberos for the first time.

For more information, see the section on Kerberos security services in the *Data ONTAP 7-Mode File Access and Protocols Management Guide*.

Preparing for the upgrade

Before installing the latest Data ONTAP release on your storage system, you need to verify information and complete some tasks.

Steps

1. Verify that your system meets the minimum requirements.
2. Verify that you have resolved any upgrade issues.
3. Ensure that you have a current Snapshot copy of the root volume of any system being upgraded.

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

4. Verify whether you need to update disk or disk shelf firmware.

Note: You should ensure that any required disk firmware and disk shelf firmware updates have completed before beginning a nondisruptive upgrade. In particular, you should upgrade disk firmware at least one day before beginning a nondisruptive Data ONTAP upgrade, because automatic background disk firmware updates can take a long time in large-capacity systems.

5. If you have storage systems in high-availability configurations, verify that they are correctly configured using the HA Configuration Checker.
6. If you are running SnapMirror, identify storage systems with destination volumes and upgrade them before upgrading storage systems with source volumes.
7. If you are running MetroCluster systems, verify that all MetroCluster components are compatible with the target release.

For more information, see your MetroCluster documentation and the MetroCluster Compatibility Matrix. If you are running MetroCluster on a V-Series system, see also the *V-Series Support Matrix*.

8. Check whether you need to perform one or both of the procedures described in the following table.

If...	Then complete this procedure...
You are running CIFS on the storage system and are using a Windows NT 4.0 domain controller for authentication	Verify that the storage system has a domain account
You are running CIFS on the storage system and are using a Windows 2000 domain controller for authentication	Enable DNS with Windows 2000 name server addresses

9. If you are using the nondisruptive upgrade method, ensure that your systems meet the requirements.

10. If you are upgrading from a release earlier than Data ONTAP 7.3, ensure that there is adequate free space in your aggregates.
11. If you have configured IPv6, ensure that you have comparable IPv4 connectivity before upgrading.

Note: IPv6 is not supported in the Data ONTAP 8.0 release family.

Next topics

[Verifying system requirements](#) on page 44

[Enabling DNS with Windows 2000 name server addresses](#) on page 47

[Verifying that you have a domain account](#) on page 47

[HA Configuration Checker](#) on page 47

[Preparing for nondisruptive upgrades](#) on page 48

[Preparing for nondisruptive upgrades on systems with VMware ESX server hosts](#) on page 50

[Determining system capacity and space guarantees before upgrading to Data ONTAP 7.3 or later](#) on page 51

[Using the aggrSpaceCheck tool to prepare your upgrade to Data ONTAP 7.3 or later](#) on page 52

[Reconfiguring IPv4 before upgrading](#) on page 53

Related concepts

[Evaluating upgrade issues](#) on page 32

[HA Configuration Checker](#) on page 47

[Why you must plan for SnapMirror upgrades](#) on page 25

Related tasks

[Detecting outdated disk firmware](#) on page 95

[Determining the required firmware for your disk shelves](#) on page 46

Related information

[MetroCluster Compatibility Matrix -- now.netapp.com/NOW/knowledge/docs/olio/guides/metrocluster_compatibility/](http://now.netapp.com/NOW/knowledge/docs/olio/guides/metrocluster_compatibility/)

[V-Series Support Matrix -- now.netapp.com/NOW/knowledge/docs/V-Series/supportmatrix/V-Series_SupportMatrix.pdf](http://now.netapp.com/NOW/knowledge/docs/V-Series/supportmatrix/V-Series_SupportMatrix.pdf)

Verifying system requirements

Before you upgrade, you must make sure your system meets the minimum requirements.

Next topics

[Ensuring that your system supports the target Data ONTAP release](#) on page 45

Ensuring that there is adequate free space in every volume containing LUNs on page 45

Deduplication upgrade requirements on page 46

Determining the required firmware for your disks on page 46

Determining the required firmware for your disk shelves on page 46

Ensuring that your system supports the target Data ONTAP release

You can check the available Data ONTAP releases on the NOW site to determine if your system supports the target Data ONTAP release.

Steps

1. Use a Web browser to go to the NOW site at now.netapp.com.
2. Click **Software** in the Download section.
3. In the **Select Platform** list box in the Data ONTAP product row, select your storage system type.
4. Click **Go**.

Result

You see a list of the releases of Data ONTAP supported by your storage system platform. If the target release is listed, you can upgrade to it.

Ensuring that there is adequate free space in every volume containing LUNs

Before upgrading a storage system in a SAN environment, you must ensure that every volume containing LUNs has available at least 1 MB of free space. The space is needed to accommodate changes in the on-disk data structures used by the new version of Data ONTAP.

About this task

"LUNs" in this context refers to the LUNs that Data ONTAP serves to clients, not to the array LUNs used for storage on a storage array.

Steps

1. Check free space in a volume containing LUNs by entering the following command at the storage system command line:

```
df
```
2. If the volume does not have at least 1 MB (1024 KB) of free space, create free space in the full volume either by deleting unnecessary data or by growing the size of the volume.

Deduplication upgrade requirements

When you upgrade to the Data ONTAP 8.0 release family, you must ensure that there is at least 1 MB of free space in each deduplicated volume. Otherwise, the deduplication metadata is not upgraded and deduplication is disabled on those volumes.

If you receive a deduplication failure message, you should add space to the FlexVol volume and run deduplication again.

Determining the required firmware for your disks

By viewing the latest required firmware revisions for Fibre Channel and SAS disk drives on the NOW site, you can determine if you need to update the disk firmware for your system.

Steps

1. Use a Web browser to go to the NOW site at now.netapp.com.
2. Click **Software** in the Download section.
3. Select **Firmware > Disk Drive & Firmware Matrix**.
4. When the Disk Drive & Firmware Matrix appears, click the link to the firmware revision needed for drives attached to your system.

Result

A page is displayed that includes installation procedures, other information about the disk firmware, and links to download images.

Determining the required firmware for your disk shelves

By viewing the latest required firmware revisions for disk shelves on the NOW site, you can determine if you need to update the disk shelf firmware for your system.

Steps

1. Use a Web browser to go to the NOW site at now.netapp.com.
2. Click **Software** in the Download section.
3. Select **Firmware > Disk Shelf Firmware**.
4. When the Disk Shelf Firmware page appears, click the link to the firmware revision needed for disk shelves attached to your system.

Result

A page is displayed that includes installation procedures, other information about the shelf firmware, and links to download images.

Enabling DNS with Windows 2000 name server addresses

If you are running CIFS on the storage system and are using a Windows 2000 domain controller for authentication, then before upgrading, you need to enable DNS with Windows 2000 name server addresses.

Steps

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 ip_address
```

Example

```
nameserver 192.9.200.10
```

2. Save the file.
3. Enter the following command at the storage system console to enable DNS:

```
options dns.enable on
```

Verifying that you have a domain account

If you are running CIFS and using a Windows NT 4.0 domain controller for authentication, you need to verify that your storage system has a domain account.

Step

1. From the storage system's console, enter the following command:

```
cifs domaininfo
```

Data ONTAP displays the storage system's domain information.

HA Configuration Checker

Before upgrading your HA configuration, you must verify that it is properly configured. You can use the HA Configuration Checker to identify and resolve any high-availability configuration issues before continuing with the upgrade. The utility is available on the NOW site.

The HA Configuration Checker (formerly the Cluster Configuration Checker) is a utility that detects errors in the configuration of a pair of high-availability storage systems.

For more information about using this utility for high-availability configuration management, see the *Data ONTAP 7-Mode High-Availability Configuration Guide*.

Related information

[HA Configuration Checker -- now.netapp.com/NOW/download/tools/cf_config_check/](https://now.netapp.com/NOW/download/tools/cf_config_check/)

Preparing for nondisruptive upgrades

You must complete certain steps to ensure a successful nondisruptive upgrade procedure. Configurations that are eligible for nondisruptive upgrades must meet certain protocol and availability requirements.

About this task

Ensure that you understand these requirements before you use the nondisruptive method.

Note: Be sure to use the Upgrade Advisor tool (if it is available in your environment) to help you determine nondisruptive upgrade requirements.

Steps

1. Ensure that your HA pair is optimally configured and functioning correctly.

The system clocks on both partner systems should be synchronized with a time server. A discrepancy in system time between the partner systems could cause problems with the upgrade.

You can verify that your HA pair is properly configured by running the HA Configuration Checker.

2. Ensure that your clients are optimally configured and functioning correctly.

Check service protocols and configure client timeout settings to ensure availability meets requirements for a nondisruptive upgrade.

3. Verify that all components of your SAN configuration are compatible with the upgraded Data ONTAP release by consulting *NetApp Interoperability Matrix* on the NOW site.

4. If the automatic giveback option, `cf.giveback.auto.enable`, is set to `on`, disable automatic giveback by entering the following command on one of your storage systems in the high-availability configuration:

```
options cf.giveback.auto.enable off
```

After the upgrade procedure, you can reset this option to `on` (if desired).

5. Ensure that you have no failed disks on either node.

If either node has failed disks, giveback might fail. To avoid this issue, remove any failed disks before entering the `cf giveback` command.

6. Remove any old core files from the `/etc/crash` directory.

For more information about managing the contents of the `/etc/crash` directory and deleting old core files, see the `savecore(1)` man page.

7. If you need disk firmware updates in addition to the Data ONTAP upgrade, ensure that all disks on your system are in RAID-DP or mirrored RAID4 aggregates.

Disk firmware updates take place automatically in the background when RAID-DP protection is configured. Services and data continue to be available during the disk firmware update.

Note: RAID4 volumes can be upgraded nondisruptively (temporarily or permanently) to RAID-DP to automatically enable the background firmware update capability.

8. If you are upgrading to this Data ONTAP release from an earlier release family, ensure that your disk firmware and disk shelf firmware are current. If they are not, you must update to the latest disk firmware and disk shelf firmware before starting the nondisruptive upgrade procedure.
9. If you use deduplication technology, ensure that your system includes no more than 300 deduplicated volumes and that no deduplication operations are active during the Data ONTAP upgrade.
10. If you use SnapMirror technology, ensure that SnapMirror is suspended and no SnapMirror operations are in process while upgrading Data ONTAP.
11. If you are planning to perform a nondisruptive upgrade on a system that does not send AutoSupport messages, you should nonetheless trigger AutoSupport notifications using the `autosupport.doit` option at the beginning and end of the upgrade.

These notifications allow you to preserve a local copy of information about the state of your system before the upgrade.

Related concepts

[HA Configuration Checker](#) on page 47

[Optimal service availability during upgrades](#) on page 131

[Disk firmware updates](#) on page 92

[Disk shelf firmware updates](#) on page 97

Related tasks

[Using the Upgrade Advisor to plan your upgrade](#) on page 19

Preparing for nondisruptive upgrades on systems with VMware ESX server hosts

Before performing a nondisruptive upgrade on storage systems exporting data over NFS to VMware ESX server hosts, verify that your client's NAS components are correctly configured, to ensure service availability for VMware guest operating systems during the upgrade.

About this task

These steps must be performed from the ESX server or guest operating systems, not from the storage system.

Steps

1. Increase the NFS datastore's heartbeat time on the VMware ESX server.

The following parameters should be set to the recommended values:

Parameter	Value
NFS.HeartbeatFrequency	12
NFS.HeartbeatMaxFailures	10

For more information about setting ESX server parameters, see the ESX documentation.

2. Set the SCSI Disk timeout value on all guest operating systems to 190 seconds.

You can obtain scripts to set the recommended SCSI disk settings in the guest operating systems for use with VMware ESX 3.5 and storage systems running Data ONTAP. When downloaded and run on the guest operating systems, the scripts create and modify the necessary files for each guest operating system type. Using the scripts ensures that the correct timeout settings are used in the guest operating systems to achieve maximum I/O resiliency when the guest operating systems are connected to storage systems.

For more information about obtaining and running the scripts, see the knowledgebase article *VMware ESX Guest OS I/O Timeout Settings for NetApp Storage Systems* on the NOW site.

3. Align the file systems that use virtual machine disk format (VMDK) on Windows with the storage systems' WAFL file system.

This step is optional but recommended for best performance.

Virtual machines store their data on virtual disks. As with physical disks, these disks are formatted with a file system. When formatting a virtual disk, the file systems with VMDK format, the datastore, and the storage array should be in proper alignment. Misalignment of the virtual machine's file system can result in degraded performance.

When aligning the partitions of virtual disks for use with storage systems, the starting partition offset value must be divisible by 4,096. The recommended starting offset value for Windows

2000, 2003, and XP operating systems is 32,768. Windows 2008 and Vista default at 1,048,576; that value does not require any adjustments.

For more information about aligning virtual disks and WAFL file systems, see "Virtual Machine Partition Alignment" in the Technical Report *NetApp and VMware Virtual Infrastructure 3, Storage Best Practices*.

Related information

NetApp and VMware Virtual Infrastructure 3, Storage Best Practices: media.netapp.com/documents/tr-3428.pdf

VMware ESX Guest OS I/O Timeout Settings for NetApp Storage Systems: now.netapp.com/Knowledgebase/solutionarea.asp?id=kb41511

Determining system capacity and space guarantees before upgrading to Data ONTAP 7.3 or later

If you suspect that your system has almost used all of its free space, or if you use thin provisioning, you should check the amount of space in use by each aggregate. If any aggregate is 97 percent full or more, *do not* proceed with the upgrade until you have used the Upgrade Advisor or aggrSpaceCheck tools to determine your system capacity and plan your upgrade.

Steps

1. Check your system's capacity by entering the following command:

```
df -A
```

If the capacity field shows...	Then...
96% or less for all aggregates	You can proceed with your upgrade to Data ONTAP 7.3; no further action is required.
97% or more for any aggregate	Continue with Step 2.

2. Use a Web browser to go to the NOW Web site and select the appropriate upgrade tool for your environment.

If your system is...	Then continue to plan your upgrade with the...
Configured to send AutoSupport messages to NetApp	Upgrade Advisor tool.
Not configured to send AutoSupport messages to NetApp	aggrSpaceCheck tool.

These tools can assess the free space requirements for your system in Data ONTAP 7.3 and later releases. If you do not have sufficient free space, the tools will recommend a course of action to ensure a successful upgrade.

Note: The Upgrade Advisor can also assess other upgrade requirements and recommend solutions. Using it is the recommended upgrade methodology.

After you finish

After using these tools and completing the upgrade, make sure that your space guarantees are configured according to your requirements.

Related tasks

[Using the Upgrade Advisor to plan your upgrade](#) on page 19

[Using the aggrSpaceCheck tool to prepare your upgrade to Data ONTAP 7.3 or later](#) on page 52

Using the aggrSpaceCheck tool to prepare your upgrade to Data ONTAP 7.3 or later

You must use the aggrSpaceCheck tool if any aggregate on your storage system is 97 percent full or more, and if you cannot use the Upgrade Advisor tool.

Before you begin

If your current system capacity is 96 percent or less for all aggregates, you do not need to complete this procedure. You can proceed with your upgrade to Data ONTAP 7.3 and later releases.

To use the aggrSpaceCheck tool, you must have the following:

- A Windows or UNIX client system with RSH enabled
- RSH configured on your storage system(s)

For information about configuring RSH, see the *Data ONTAP 7-Mode System Administration Guide*.

- Access to the NOW site
- Access to the storage system being upgraded
- Root user privileges

About this task

The aggrSpaceCheck tool is a utility that runs on the administration host client system. It is available for download from the ToolChest area on the NOW site. When installed on the client system, it connects to the storage system using the RSH protocol and checks whether there is enough free space to enable Data ONTAP 7.3. It does so by executing several Data ONTAP commands, parsing the result, and performing calculations to assess space requirements. The results and recommended actions are displayed immediately.

Steps

1. From a Web browser, log in to the ToolChest page on NOW at the following URL:

<https://now.netapp.com/eservice/toolchest>

2. Enter one of the following commands, depending on your client system.

If you have a... Enter the following command...

Windows client `aggrSpaceCheck [-user user_name] -filer system_name`

UNIX client `perl aggrSpaceCheck.pl [-user user_name] -filer system_name`

Example

To connect to a system called server1, enter the following command from a Windows client:

```
aggrSpaceCheck -filer server1
```

To connect to a system called server1 as user sysadmin, enter the following command from a Windows client:

```
aggrSpaceCheck -user sysadmin -filer server1
```

To connect to a system called server1 as user root, enter the following command from a UNIX client:

```
perl aggrSpaceCheck.pl -user root -filer server1
```

For more information, see the readme.txt file that is included with the aggrSpaceCheck tool.

3. Use the recommendations displayed by the aggrSpaceCheck tool to prepare your system.

After you finish

When you have completed your preparations, proceed with the upgrade.

Related information

[aggrSpaceCheck -- now.netapp.com/eservice/toolchest](https://now.netapp.com/eservice/toolchest)

Reconfiguring IPv4 before upgrading

Before you upgrade to the Data ONTAP 8.0 release family, any configuration with only an IPv6 address must be reconfigured with an IPv4 address. In particular, you must manually reconfigure the vFiler units, CIFS, DNS servers, NIS servers, and the configuration files inside the `/etc` directory for IPv4 networking.

Steps

1. If your system includes the following configurations, complete the appropriate steps before upgrading:

If you have configured...	Then...
IPv6 addresses on any of your system's vFiler units	Reconfigure the vFiler units with IPv4 addresses. Note: Any vFiler units with IPv6 addresses cannot be reached after upgrading.
Your storage system to query DNS servers with IPv6 addresses	Reconfigure the DNS servers with IPv4 addresses by either running the <code>setup</code> command or editing the <code>/etc/resolv.conf</code> file.
Your storage system to query NIS servers with IPv6 addresses	Reconfigure the NIS servers with IPv4 addresses by running either the <code>setup</code> command or the <code>options nis.servers</code> command. You can also edit the <code>/etc/hosts</code> file to replace all IPv6 addresses with IPv4 addresses.

Note: For configuring the DNS and NIS servers from the vfiler context, follow the steps in the preceding table from the vfiler context.

2. To remove the IPv6 addresses from the `/etc/exports` file after upgrading, you can edit the file manually and remove the IPv6 addresses.

This step is optional.

You can use the `exportfs -w` command to write the export rules that are stored in the memory to the `/etc/exports` file. This command removes all IPv6 addresses from the `/etc/exports` file. The `/etc/exports` file with IPv6 addresses is backed up to the `/etc/exports.bak2` file.

3. From a workstation that has access to your storage system's root volume, open the `/etc/hosts` and `/etc/rc` files by using a text editor and replace all IPv6 address configuration in these files with IPv4 addresses.
4. Reboot the storage system.
5. Verify the IPv4 connectivity before upgrading.

Obtaining Data ONTAP software images

You must copy a software image from the NOW site to your storage system using UNIX or Windows client connections. Alternatively, you can copy software images to an HTTP server on your network and then storage systems can access the images using the `software` command.

To upgrade the storage system to the latest release of Data ONTAP, you need access to software images, software version information, and the latest firmware for your storage system model are available on the NOW site. Note the following important information:

- Software images are specific to storage system models.
Be sure to obtain the correct image for your system.
- Software images include the latest version of system firmware that was available when a given version of Data ONTAP was released.

Attention: Beginning with Data ONTAP 8.0, .exe images are no longer used for Data ONTAP software upgrades. You must use one of the following image types, depending on the upgrade you are performing:

- .zip images, for upgrades from an earlier release family to Data ONTAP 8.0
- .tgz images, for upgrades from any Data ONTAP 8.0 release to a later release

After you have upgraded to Data ONTAP 8.0 or later, you can only use .tgz images for further upgrades.

Next topics

[Obtaining images for HTTP servers](#) on page 56

[Obtaining images for UNIX clients](#) on page 57

[Obtaining images for Windows clients](#) on page 59

[Managing files in the /etc/software directory](#) on page 61

Related information

[Download Software -- now.netapp.com/NOW/cgi-bin/software](#)

[System Firmware + Diagnostics Download -- now.netapp.com/NOW/cgi-bin/fw](#)

Obtaining images for HTTP servers

If you have an HTTP server that is accessible to your storage system, you can copy Data ONTAP software images to the HTTP server and use the `software` command to download and install Data ONTAP software images to your storage system.

Note: You can also use HTTPS connections when SecureAdmin is installed and enabled on the storage system.

When you use an HTTP server to provide Data ONTAP software images, 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 copy the Data ONTAP system files to both single systems and storage systems in a high-availability configuration.

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

Next topics

[Copying the software image to the HTTP server](#) on page 56

[Copying software images from the HTTP server without installing the images](#) on page 56

Related concepts

[Installing Data ONTAP software images](#) on page 63

Copying the software image to the HTTP server

You must copy the software image file to the HTTP server. This task prepares the HTTP server to serve software images to storage systems in your environment.

Step

1. Copy the software image (for example, `80_setup_i.tgz`) from the NOW site or another system to the directory on the HTTP server from which the file will be served.

Copying software images from the HTTP server without installing the images

You can copy software images to your storage system without immediately installing them. You might do this, for instance, if you want to perform the installation at a later time.

Step

1. Enter the following command from the storage system console:

```
software get url -f filename
```

url is the HTTP location from which you want to copy the Data ONTAP software images.

Use the following URL syntax if you need to specify a user name, password, host, and port to access files on the HTTP server using Basic Access Authentication (RFC2617):

`http://username:password@host:port/path`

Use the `-f` flag to overwrite an existing software file of the same name in the storage system's `/etc/software` 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`.

filename 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 `/etc/software` directory on the storage system.

Example

In the following example, the `software get` command uses a new destination file name:

```
software get http://www.example.com/downloads/x86-64/80_setup_i.tgz  
80_mailboxes_i.tgz
```

You see a message similar to the following:

```
software: copying to /etc/software/80_mailboxes_i.tgz  
software: 100% file read from location.  
software: /etc/software/80_mailboxes_i.tgz has been copied.
```

Obtaining images for UNIX clients

If you are using a UNIX client to copy a Data ONTAP software image to your storage system, you need access to both the storage system's console and the system's upgrade host. If the upgrade host does not have a Web connection, you must also have access to a client system that can reach the NOW site.

Next topics

[Mounting the storage system on your client](#) on page 58

[Obtaining software images](#) on page 58

Related concepts

[Upgrade host requirements](#) on page 23

[Installing Data ONTAP software images](#) on page 63

Mounting the storage system on your client

Before you copy a software image to your storage system, you must mount the system on your UNIX upgrade host.

Steps

1. As root user, mount the storage system's root file system to the client's `/mnt` directory, using the following command:

```
mount system:/vol/vol0 /mnt
```

system is the name of the storage system.

`/mnt` is the directory on the client where you want to mount the storage system's root file system.

2. Change to the `/mnt` directory using the following command on your UNIX client console:

```
cd /mnt
```

`/mnt` is the directory on the client where you mounted the storage system's root file system.

3. To acquire Data ONTAP files, download the Data ONTAP files using a Web browser from the NOW site.

Obtaining software images

You can use a Web browser to copy the software image from the NOW site to a UNIX client.

About this task

You can copy the software image directly to your upgrade host. If your upgrade host does not have Web access, you can copy the software image to portable storage media attached to a different client, then copy the image from portable storage to the upgrade host.

Steps

1. Use a Web browser to log in to the NOW site.
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. After you have chosen the software image that corresponds to your platform, complete one of the following actions, depending on your Web environment.

If you are connecting to the NOW site from... Then...

An upgrade host	Save the image to the <code>.../etc/software</code> directory on the mountpoint that you chose when you mounted the storage system on your client.
Another UNIX client	<ol style="list-style-type: none"> a. Save the image to portable storage media. b. Connect the portable storage media to your upgrade host. c. Copy the image to the <code>.../etc/software</code> directory on the mountpoint that you chose when you mounted the storage system on your client.

8. Continue with the installation procedures.

Obtaining images for Windows clients

If you are using a Windows client to copy a Data ONTAP software image to your storage system, you need access to both the storage system's console and the system's upgrade host. If the upgrade host does not have a Web connection, you must also have access to a client system that can reach the NOW site.

Next topics

[Mapping the storage system to a drive](#) on page 59

[Obtaining software images](#) on page 60

Related concepts

[Upgrade host requirements](#) on page 23

[Installing Data ONTAP software images](#) on page 63

Mapping the storage system to a drive

Before you copy a software image to your storage system, you must map the root directory of the system to your Windows upgrade host.

Before you begin

You should make sure that the CIFS service is running and that the Administrator user is defined in CIFS as having authority to access the C\$ directory.

Steps

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.

Note: On some computers, firewall software might not permit you to map a drive to the C\$ directory of a storage system. To complete this procedure, disable the firewall until you no longer need access to the storage system through your laptop.

3. Copy the software image from the NOW site.

Obtaining software images

You can use a Web browser to copy the software image from the NOW site to a Windows client.

About this task

You can copy the software image directly to your upgrade host. If your upgrade host does not have Web access, you can copy the software image to portable storage media attached to a different client, then copy the image from portable storage to the upgrade host.

Steps

1. Use a Web browser to log in to the NOW site.
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. After you have chosen the software image that corresponds to your platform, complete one of the following actions, depending on your Web environment.

If you are connecting to the NOW site from...	Then...
An upgrade host	Save the image to the <code>\etc\software</code> directory on the mountpoint that you chose previously, when you mounted the storage system on your client.
Another Windows client	<ol style="list-style-type: none"> a. Save the image to portable storage media. b. Connect the portable storage media to your upgrade host. c. Copy the image to the <code>\etc\software</code> directory on the mountpoint that you chose previously, when you mounted the storage system on your client.

8. Continue with the installation procedures.

Managing files in the `/etc/software` directory

After you have copied Data ONTAP system files to the `/etc/software` directory on your storage system, you can manage them from the storage system console with the `software` command.

If you want to...	Then use the following command...
List the contents of the <code>/etc/software</code> directory	<code>software list</code>
Delete files from the <code>/etc/software</code> directory	<code>software delete</code>

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

Installing Data ONTAP software images

You should use the `software update` command to extract and install the system files on a storage system.

You can use the `software update` command to install a software image you have already copied to your storage system, or to copy and install the image from an HTTP server.

You must know the location of and have access to the software image. The `software update` command requires one of the following as an argument:

- The name of the software image you copied to the `/etc/software` directory
- The URL of the HTTP server that you configured to serve software images

The `software update` command allows you to perform several operations at one time. For example, if you use an HTTP server to distribute software images, you can copy an image from the HTTP server, extract and install the system files, download the files to the boot device, and reboot your system with one command.

For more information about the `software update` command and its options, see the `software(1)` man page.

Note: Beginning with Data ONTAP 8.0, the following processes are no longer supported for extracting and installing Data ONTAP software images:

- Using the `tar` command from UNIX clients
- Using the `setup.exe` file and WinZip from Windows clients

For the upgrade to Data ONTAP 8.0 and later releases, `.exe` images are no longer available. You must use one of the following image types depending on the upgrade you are performing:

- `.zip` images, for upgrades from an earlier release family to Data ONTAP 8.0
- `.tgz` images, for upgrades from any Data ONTAP 8.0 release to a later release

After you have upgraded to Data ONTAP 8.0 or later, `.tgz` images are the only image type you can use for further upgrades.

Next topics

[*Installing software images from an HTTP server*](#) on page 64

[*Installing software images from the `/etc/software` directory*](#) on page 67

Installing software images from an HTTP server

To complete this procedure, you must know the URL of an HTTP server in your environment that is configured to serve software images.

Step

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

```
software update url options
```

- *url* is the URL of the HTTP server and subdirectory.
- *options* is one or more of the following:
 - The `-d` option prevents the `download` command from being run automatically after the system files are installed.
 - The `-f` option overwrites the existing image in the `/etc/software` directory.
 - The `-r` option prevents the system from rebooting automatically after the `download` command has finished (default).
 - The `-R` option causes the system to reboot automatically after the `download` command has finished.

Attention: Beginning in Data ONTAP 8.0.1, the `software update` options have changed; the `-r` option (no automatic reboot) is the default, and the `-R` option must be specified to override the `-r` option.

However, if you are upgrading from any release earlier than Data ONTAP 8.0.1, you must include the `-r` option to prevent automatic reboot if you are performing a nondisruptive upgrade or if you are upgrading firmware.

For more information, see the `software(1)` man page for the Data ONTAP version currently running on your system.

Example

If you are running Data ONTAP...	And you want to...	Then you can enter...
8.0.1 or later, or 7.3.5 or later	Copy and install the image from your HTTP server	<pre>software update http:// www.example.com/ downloads/x86-64/ my_80_setup_i.tgz -d</pre>
	Copy from your HTTP server and overwrite an existing image	<pre>software update http:// www.example.com/ downloads/x86-64/ my_80_setup_i.tgz -d -f</pre>
	Copy and install the image from your HTTP server, then download the new system files to the boot device immediately after installing them	<pre>software update http:// www.example.com/ downloads/x86-64/ my_80_setup_i.tgz</pre>
	Copy and install the image from your HTTP server to a single system, then download the new system files and reboot immediately	<pre>software update http:// www.example.com/ downloads/x86-64/ my_80_setup_i.tgz -R</pre>

If you are running Data ONTAP...	And you want to...	Then you can enter...
8.0, or 7.3.4 or earlier	Copy and install the image from your HTTP server	<code>software update http:// www.example.com/ downloads/x86-64/ my_80_setup_i.tgz -d - r</code>
	Copy from your HTTP server and overwrite an existing image	<code>software update http:// www.example.com/ downloads/x86-64/ my_80_setup_i.tgz -d - r -f</code>
	Copy and install the image from your HTTP server, then download the new system files to the boot device immediately after installing them	<code>software update http:// www.example.com/ downloads/x86-64/ my_80_setup_i.tgz -r</code>
	Copy and install the image from your HTTP server to a single system, then download the new system files and reboot immediately	<code>software update http:// www.example.com/ downloads/x86-64/ my_80_setup_i.tgz</code>

When you use the `software update` command without the options, a message similar to the following appears on your storage system console:

```
software: You can cancel this operation by hitting Ctrl-C in the next 6
seconds.
software: Depending on system load, it might take many minutes
software: to complete this operation. Until it finishes, you will
software: not be able to use the console.
software: copying to <filename>
software: 100% file read from location.
software: /etc/software/<filename> has been copied.
software: installing software, this could take a few minutes...
software: Data ONTAP Package Manager Verifier 1
software: Validating metadata entries in /etc/boot/NPM_METADATA.txt
software: Checking sha1 checksum of file checksum file: /etc/boot/
NPM_FCSUM-pc.shal.asc
software: Checking sha1 file checksums in /etc/boot/NPM_FCSUM-
pc.shal.asc
software: installation of <filename> completed.
Mon Oct 2 13:26:17 PDT [filer: rc:info]: software: installation of
<filename> completed.
```

```
software: Reminder: You might need to upgrade Volume SnapMirror
destination
```

```
software: filers associated with this filer. Volume SnapMirror can not
mirror
software: if the version of ONTAP on the source filer is newer than
that on
software: the destination filer.
Mon Oct 2 13:26:17 PDT [filer: download.request:notice]
```

After you finish

Complete the installation by downloading to HA pairs or single systems.

Related concepts

[Downloading and rebooting new Data ONTAP software](#) on page 71

Installing software images from the `/etc/software` directory

To complete this procedure, the new software image must be present in the `/etc/software` directory on your storage system.

Step

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

```
software update file options
```

- *file* is the name of the software image you copied to the `/etc/software` directory.
- *options* is one or more of the following:
 - The `-d` option prevents the `download` command from being run automatically after the system files are installed.
 - The `-f` option overwrites the existing image in the `/etc/software` directory.
 - The `-r` option prevents the system from rebooting automatically after the `download` command has finished (default).
 - The `-R` option causes the system to reboot automatically after the `download` command has finished.

Attention: Beginning in Data ONTAP 8.0.1, the `software update` options have changed; the `-r` option (no automatic reboot) is the default, and the `-R` option must be specified to override the `-r` option.

However, if you are upgrading from any release earlier than Data ONTAP 8.0.1, you must include the `-r` option to prevent automatic reboot if you are performing a nondisruptive upgrade or if you are upgrading firmware.

For more information, see the `software(1)` man page for the Data ONTAP version currently running on your system.

Example

If you are running Data ONTAP...	And you want to...	Then you can enter...
8.0.1 or later, or 7.3.5 or later	Install the new system files from the <code>/etc/software</code> directory	<code>software update my_80_setup_i.tgz -d</code>
	Download the new system files to the boot device immediately after installing them	<code>software update my_80_setup_i.tgz</code>
	Copy and install the image from your HTTP server	<code>software update http://www.example.com/downloads/x86-64/my_80_setup_i.tgz</code>
	Copy from your HTTP server and overwrite an existing image	<code>software update http://www.example.com/downloads/x86-64/my_80_setup_i.tgz -f</code>
	Perform an upgrade on a single system and reboot immediately	<code>software update -R my_80_setup_i.tgz</code>
8.0, or 7.3.4 or earlier	Install the new system files from the <code>/etc/software</code> directory	<code>software update my_80_setup_i.tgz -d -r</code>
	Download the new system files to the boot device immediately after installing them	<code>software update my_80_setup_i.tgz -r</code>
	Copy and install the image from your HTTP server	<code>software update http://www.example.com/downloads/x86-64/my_80_setup_i.tgz</code>
	Copy from your HTTP server and overwrite an existing image	<code>software update http://www.example.com/downloads/x86-64/my_80_setup_i.tgz -f</code>
	Perform an upgrade on a single system and reboot immediately	<code>software update my_80_setup_i.tgz</code>

When you use the `software update` command without the options, a message similar to the following appears on your storage system console:

```
software: You can cancel this operation by hitting Ctrl-C in the next 6
seconds.
software: Depending on system load, it might take many minutes
software: to complete this operation. Until it finishes, you will
software: not be able to use the console.
software: copying to <filename>
software: 100% file read from location.
software: /etc/software/<filename> has been copied.
software: installing software, this could take a few minutes...
software: Data ONTAP Package Manager Verifier 1
software: Validating metadata entries in /etc/boot/NPM_METADATA.txt
software: Checking sha1 checksum of file checksum file: /etc/boot/
NPM_FCSUM-pc.shal.asc
software: Checking sha1 file checksums in /etc/boot/NPM_FCSUM-
pc.shal.asc
software: installation of <filename> completed.
Mon Oct 2 13:26:17 PDT [filer: rc:info]: software: installation of
<filename> completed.
```

```
software: Reminder: You might need to upgrade Volume SnapMirror
destination
software: filers associated with this filer. Volume SnapMirror can not
mirror
software: if the version of ONTAP on the source filer is newer than
that on
software: the destination filer.
Mon Oct 2 13:26:17 PDT [filer: download.request:notice]
```

After you finish

Complete the installation by downloading to HA pairs or single systems.

Related concepts

[Downloading and rebooting new Data ONTAP software](#) on page 71

Downloading and rebooting new Data ONTAP software

The upgrade method you use depends on the kind of upgrade.

If you are upgrading systems in a SnapMirror environment, you must also follow these instructions:

- Upgrade them in the correct order.
- Suspend SnapMirror operations before performing a nondisruptive upgrade.

Next topics

[Upgrading in a SnapMirror environment](#) on page 71

[Upgrading nondisruptively in a SnapMirror environment](#) on page 72

[Upgrading HA configurations from an earlier release family nondisruptively](#) on page 73

[Upgrading HA configurations within a release family nondisruptively](#) on page 78

[Upgrading HA configurations using the standard method](#) on page 81

[Upgrading single systems](#) on page 84

Related concepts

[Release family upgrade requirements](#) on page 26

[Standard upgrade requirements](#) on page 32

[Nondisruptive upgrade requirements](#) on page 28

Upgrading in a SnapMirror environment

If you need to upgrade Data ONTAP on a system that uses SnapMirror for volume replication, you must upgrade systems with destination volumes *before* you upgrade systems that have source volumes.

About this task

Note: If you are upgrading nondisruptively, you must also suspend SnapMirror operations before upgrading and resume SnapMirror operations when the upgrade is finished.

SnapMirror source volumes can be replicated to single or multiple destination volumes. Replication to multiple destination volumes is also referred to as *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 systems where the source volumes reside. In addition, when you upgrade storage systems in a cascading series, you should upgrade the systems in order, beginning with the destination systems furthest logically in your topology from the source system.

Steps

1. Identify any destination volumes by entering the following command on the storage system with the source volume:

```
snapmirror destinations
```

The `snapmirror` command lists all destination volumes, including cascaded destinations.

2. Upgrade the systems that have destination volumes, beginning with the furthest system in the topology (that is, the last system in a series of cascading destinations).
3. Upgrade the system that has the source volume.

Attention: You must upgrade the 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 reenable 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 than or the same as that on the destination system.

Related tasks

[Identifying SnapMirror destination volumes](#) on page 0

Upgrading nondisruptively in a SnapMirror environment

You must suspend SnapMirror operations before performing a nondisruptive upgrade of Data ONTAP.

About this task

The requirement to suspend SnapMirror operations applies to both synchronous and asynchronous SnapMirror modes.

For more information about SnapMirror operations, see the `snapmirror(1)` man page and the *Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide*.

Steps

1. Enter the following command on both source and destination systems to disable SnapMirror operations:

```
snapmirror off
```

As an alternative, you can set the `snapmirror.enable` option to `off`.

2. For each destination volume, enter the following command to allow existing SnapMirror transfers to finish:

```
snapmirror quiesce destination
```

Example

To quiesce transfers involving the destination volume `toaster-cl1-cn:vol1`, enter the following command:

```
snapmirror quiesce toaster-cl1-cn:vol1
```

3. Complete the nondisruptive upgrade according to your upgrade plan.

4. Enter the following command to reenable SnapMirror operations:

```
snapmirror on
```

5. Enter the following command to resume existing SnapMirror transfers:

```
snapmirror resume destination
```

Upgrading HA configurations from an earlier release family nondisruptively

You can upgrade HA pairs to a new Data ONTAP release family while maintaining storage system availability. This nondisruptive upgrade method has several steps: initiating a failover operation on one system, updating the "failed" system (and if necessary, its firmware), initiating giveback, and repeating the process on the other system. When repeating the process, you must use a special `cf` command if different release families are running on the two systems in the HA pair.

Before you begin

Before initiating the nondisruptive upgrade procedure, you need to verify that you have prepared for the upgrade by completing any prerequisite procedures. You must also ensure that you have installed Data ONTAP software onto your storage system.

Steps

1. At the console of each storage system, enter the following command to verify that the HA configuration is enabled:

```
cf status
```

The `cf status` command output should be similar to the following:

```
Cluster enabled, systemA is up.
```

If the output indicates that the HA configuration is not enabled, enter the following command to enable it:

```
cf enable
```

Then verify that the HA configuration is reenabled by entering the `cf status` command.

2. Choose the following option depending on whether you have already installed new system files.

If you...	Then...
Have already installed system files	Go to the next step.
Are installing and downloading system files in the same operation	<p data-bbox="516 309 1107 333">At the console of each system, enter the following command:</p> <pre data-bbox="516 350 919 374">software update file_name -r</pre> <p data-bbox="516 395 693 420">Then go to Step 4.</p> <p data-bbox="542 437 1212 524">Note: Beginning in Data ONTAP 8.0.1, the <code>-r</code> option (no automatic reboot) is the default. However, until you are running a release that supports this option, you must continue to specify the <code>-r</code> option.</p>

When you use the `software update` command without the `-d` option, the `download` command is executed by default.

Note: Activating Data ONTAP 8.0.x software images with the `download` process takes significantly longer than in earlier releases. The `download` process for Data ONTAP 8.0.x usually takes 20 to 60 minutes.

3. At the console of each system, enter the following command to activate the new code on the storage system's boot device:

download

After some configuration reminders, the `download` command provides an acknowledgment similar to the following:

```
Tue Jun 19 10:03:22 GMT [download.request:notice]:
Operator requested download initiated
download: Downloading boot device
.....
download: Downloading boot device (Service Area)
```

Then a message similar to the following appears:

```
Tues Jun 19 10:11:51 GMT [download.requestDone:notice]:
Operator requested download completed
```

Note: The storage system console is unavailable until the `download` procedure is complete.

4. Choose the following option that describes your system configuration.

If CIFS...	Then...
Is not in use in system A	Go to the next step.
Is in use in system A	<p data-bbox="462 1473 758 1498">Enter the following command:</p> <pre data-bbox="462 1515 758 1539">cifs terminate -t nn</pre> <p data-bbox="462 1557 1233 1616"><code>nn</code> is a notification period (in minutes) appropriate for your clients after which CIFS services are terminated. After that period of time, proceed to the next step.</p>

- At the console of system B, enter the following command:

```
cf takeover
```

This command causes system A to shut down gracefully and leaves system B in takeover mode.

- To display the LOADER boot prompt at the system A console, press Ctrl-C at the system A console when instructed after the boot sequence starts.

You can also display the LOADER prompt by pressing Ctrl-C at the system A console when the "Waiting for giveback" message appears at the console of system A. When prompted to halt the node rather than wait, enter **y**.

- After halting the node, check the Boot Loader messages for a warning similar to the following: Warning: The CompactFlash contains newer firmware image (1.6.0). Please run 'update_flash' at Loader prompt to update your system firmware (1.5X3).

If you...	Then ...
Do not see this warning.	BIOS firmware is updated automatically if needed; go to Step 11.
See this warning.	You must update BIOS firmware manually; go to the next step.

After the new BIOS system firmware is installed, future system firmware updates take place automatically.

Attention: The version of BIOS system firmware that ships with Data ONTAP 8.0 and higher is the minimum requirement for running this Data ONTAP release. If your system is running a lower level of system firmware, you will not be able to boot Data ONTAP 8.0 and the upgrade will fail.

- At the boot prompt, enter the following command to reset the system:

```
bye
```

- Display the LOADER boot prompt again at the system A console by repeating Step 6.
- Enter the following command:

```
update_flash
```

The system updates the firmware, displays several status messages, and displays the boot prompt.

- Enter the following command to reboot the system using the new firmware and software:

```
bye
```

- Choose the option that describes your configuration.

If FCP or iSCSI...	Then when the "Waiting for giveback" message appears on the console of system A...
Is not in use in system A	Enter the following command at the console of system B: cf giveback

If FCP or iSCSI...	Then when the "Waiting for giveback" message appears on the console of system A...
Is in use in system A	Wait for at least eight minutes to allow host multipathing software to stabilize, then enter the following command at the console of system B: cf giveback

Attention: The `cf giveback` command can fail because of open client sessions (such as CIFS sessions), long-running operations, or operations that cannot be restarted (such as tape backup or SyncMirror resynchronization). If the `cf giveback` command fails, terminate any CIFS session or long-running operations gracefully (because the `-f` option will immediately terminate any CIFS sessions or long-running operations) and then enter the following command (with the `-f` option):

```
cf giveback -f
```

For more information about the behavior of the `-f` option, see the `cf(1)` man page.

The command causes system A to reboot with the new system configuration—a Data ONTAP version and any new system firmware and hardware changes—and resume normal operation as a high-availability partner.

Note: At this point in the upgrade procedure—system A is running the new Data ONTAP version and system B is running an earlier Data ONTAP release family—the systems are in a state of "version mismatch." This means that normal high-availability functions such as NVRAM mirroring and automatic takeover are not in effect. You might see error messages indicating version mismatch and mailbox format problems. This is expected behavior; it represents a temporary state in a major nondisruptive upgrade and is not harmful.

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.

There might be several reboots if component firmware needs to be updated. These interim reboots do not affect the nondisruptive upgrade; the final reboot returns the system to high-availability status.

13. Choose the following option that describes your configuration.

If CIFS...	Then...
Is not in use in system B	Go to the next step.
Is in use in system B	Enter the following command: cifs terminate -t nn <i>nn</i> is a notification period (in minutes) appropriate for your clients after which CIFS services are terminated. After that period of time, proceed to the next step.

14. At the console of system A, enter the following command:

cf takeover -n

You see output similar to the following:

```
Waiting for partner to be cleanly shutdown using the
'halt' command
Press Ctrl-C to abort wait...
```

Note: The `-n` flag of the `cf takeover` command should only be used for major nondisruptive upgrades. If run during a minor nondisruptive upgrade or a non-upgrade takeover, it will generate an error and the command will terminate.

15. At the console of system B, enter the following command:

halt

This command causes system B to shut down cleanly, flushing file-system information in memory to disk.

16. After halting the node, check the Boot Loader messages for a warning similar to the following:
Warning: The CompactFlash contains newer firmware image (1.6.0). Please run 'update_flash' at Loader prompt to update your system firmware (1.5X3).

If...	Then...
You do not see this warning	BIOS firmware is updated automatically if needed; go to Step 20.
You see this warning	You must update BIOS firmware manually; go to the next step.

After the new BIOS system firmware is installed, future system firmware updates take place automatically.

Attention: The version of BIOS system firmware that ships with Data ONTAP 8.0 and higher is the minimum requirement for running this Data ONTAP release. If your system is running a lower level of system firmware, you will not be able to boot Data ONTAP 8.0 and the upgrade will fail.

17. At the boot prompt, enter the following command to reset the system:

bye

18. To display the LOADER boot prompt at the system B console, press Ctrl-C at the system B console when instructed after the boot sequence starts.

You can also display the LOADER prompt by pressing Ctrl-C at the system A console when the "Waiting for giveback" message appears at the console of system B. When prompted to halt the node rather than wait, enter **y**.

19. Enter the following command:

update_flash

The system updates the firmware, displays several status messages, and displays the boot prompt.

20. 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
```

21. Choose the option that describes your configuration.

If FCP or iSCSI...	Then when the "Waiting for giveback" message appears on the console of system B...
Is not in use in system B	Enter the following command at the console of system A: <code>cf giveback</code>
Is in use in system B	Wait for at least eight minutes to allow host multipathing software to stabilize, then enter the following command at the console of system A: <code>cf giveback</code>

Attention: The `cf giveback` command can fail because of open client sessions (such as CIFS sessions), long-running operations, or operations that cannot be restarted (such as tape backup or SyncMirror resynchronization). If the `cf giveback` command fails, terminate any CIFS session or long-running operations gracefully (because the `-f` option will immediately terminate any CIFS sessions or long-running operations) and then enter the following command (with the `-f` option):

```
cf giveback -f
```

For more information about the behavior of the `-f` option, see the `cf(1)` man page.

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 high-availability partner.

Note: There might be several reboots if component firmware needs to be updated. These interim reboots will not affect the nondisruptive upgrade; the final reboot returns the system to high-availability status.

When the final reboot is finished, the two high-availability nodes are running the same Data ONTAP version.

Upgrading HA configurations within a release family nondisruptively

You can upgrade HA pairs within a Data ONTAP release family while maintaining storage system availability. This nondisruptive upgrade method has several steps: initiating a failover operation on

one system, updating the "failed" system (and if necessary, its firmware), initiating giveback, and repeating the process on the other system.

Before you begin

Before initiating the nondisruptive upgrade procedure, you need to prepare for the upgrade by completing any prerequisite procedures. You must also ensure that you installed the Data ONTAP software onto your storage system.

Steps

1. At the console of each storage system, enter the following command to verify that the HA configuration is enabled:

cf status

The `cf status` command output should be similar to the following:

Cluster enabled, systemA is up.

If the output indicates that the HA configuration is not enabled, enter the following command to enable it:

cf enable

Then verify that the HA configuration is reenabled by entering the `cf status` command.

2. Choose the following option depending on whether you have already installed new system files.

If you...	Then...
Have already installed system files	Go to the next step.
Are installing and downloading system files in the same operation	At the console of each system, enter the following command: software update file_name -r Then go to Step 4. Note: Beginning in Data ONTAP 8.0.1, the <code>-r</code> option (no automatic reboot) is the default. However, until you are running a release that supports this option, you must continue to specify the <code>-r</code> option.

When you use the `software update` command without the `-d` option, the `download` command is executed by default.

Note: Using the `download` procedure to activate Data ONTAP 8.0 software images takes significantly longer to complete than on earlier releases.

3. At the console of each system, enter the following command to activate the new code on the storage system's boot device:

download

The download command provides an acknowledgment similar to the following:

```
Tue Jun 19 10:03:22 GMT [download.request:notice]:
Operator requested download initiated
download: Downloading boot device
.....
download: Downloading boot device (Service Area)
```

Then a message similar to the following appears:

```
Tues Jun 19 10:11:51 GMT [download.requestDone:notice]:
Operator requested download completed
```

Note: The storage system console is unavailable until the download procedure is complete.

4. Choose the following option that describes your configuration.

If CIFS...	Then...
Is not in use in system A	Go to the next step.
Is in use in system A	Enter the following command: cfifs terminate -t nn <i>nn</i> is a notification period (in minutes) appropriate for your clients after which CIFS services are terminated. After that period of time proceed to Step 3.

5. At the console of system B, enter the following command:

```
cf takeover
```

This command causes system A to shut down gracefully and leaves system B in takeover mode.

6. Choose the option that describes your configuration.

If FCP or iSCSI...	Then when the "Waiting for giveback" message appears on the console of system A...
Is not in use in system A	Enter the following command at the console of system B: cf giveback
Is in use in system A	Wait for at least eight minutes to allow host multipathing software to stabilize, then enter the following command at the console of system B: cf giveback

Attention: The `cf giveback` command can fail because of open client sessions (such as CIFS sessions), long-running operations, or operations that cannot be restarted (such as tape backup or SyncMirror resynchronization). If the `cf giveback` command fails, terminate any CIFS session or long-running operations gracefully (because the `-f` option will immediately terminate any CIFS sessions or long-running operations) and then enter the following command (with the `-f` option):

```
cf giveback -f
```

For more information about the behavior of the `-f` option, see the `cf(1)` man page.

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 high-availability partner.

Note: There might be several reboots if component firmware needs to be updated. These interim reboots do not affect the nondisruptive upgrade; the final reboot returns the system to high-availability status.

7. Choose the following option that describes your configuration:

If FCP or iSCSI...	Then...
Is not in use in system A	Repeat Step 4 through Step 8 to update the partner storage system; in other words, bring down and update system B with partner A in takeover mode.
Is in use in system A	After system A resumes normal operation as a high-availability partner, wait for at least eight minutes to allow host multipathing software to stabilize. Then repeat Step 4 through Step 8 to update the partner storage system; in other words, bring down and update system B with system A in takeover mode.

8. Choose the following option that describes your configuration:

If you are upgrading from...	Then...
Data ONTAP 7.2.4 or later with AutoSupport enabled	Your nondisruptive upgrade is complete.
Any release earlier than 7.2.4, or your system is not configured to send AutoSupport messages	Trigger another AutoSupport notification by entering the following command at the console of each storage system controller: <code>options autosupport.doit finishing_NDU</code>

This notification includes a record of the system status after upgrading. It saves useful troubleshooting information in case there is a problem with the upgrade process.

Upgrading HA configurations using the standard method

If you can take HA pairs offline to update software and other components, you can use the standard upgrade method. This method has several steps: disabling the HA configuration from the console of one of the systems, updating each system (and if necessary, its firmware), and finally reenabling the HA configuration between the two systems.

Before you begin

Before initiating the standard upgrade procedure, you need to prepare for the upgrade by completing any prerequisite procedures. You must also ensure that you installed the Data ONTAP software onto the storage system.

Note: If you are upgrading a system running Data ONTAP 7.2 or later, you can use the `software update` command to complete all or part of this procedure.

Steps

1. Disable the HA configuration by entering the following command at the console of one of the storage systems:
cf disable
2. Choose the following option depending on whether you have already installed new system files:

If you...	Then...
Have already installed system files	Go to the next step.
Are installing and downloading system files in the same operation	At the console of each system, enter the following command: software update <i>file_name</i> -r Then go to Step 4. Note: Beginning in Data ONTAP 8.0.1, the <code>-r</code> option (no automatic reboot) is the default. However, until you are running a release that supports this option, you must continue to specify the <code>-r</code> option.

When you use the `software update` command without the `-d` option, the `download` command is executed by default.

Note: Activating Data ONTAP 8.0.x software images with the `download` process takes significantly longer than on earlier releases. The process for Data ONTAP 8.0.x usually finishes in 20 to 60 minutes.

3. At the console of each system, enter the following command to activate the new code on the storage system's boot device:

download

The `download` command provides an acknowledgment similar to the following:

```
Tue Jun 19 10:03:22 GMT [download.request:notice]:
Operator requested download initiated
download: Downloading boot device
.....
download: Downloading boot device (Service Area)
```

Then a message similar to the following appears:

```
Tues Jun 19 10:11:51 GMT [download.requestDone:notice]:
Operator requested download completed
```

Note: The storage system console is unavailable until the `download` procedure is complete.

4. Enter the following command at the console of system A:

```
halt
```

After the system shuts down, the LOADER prompt appears.

5. After halting the system, check the Boot Loader messages for a warning similar to the following: Warning: The CompactFlash contains newer firmware image (1.6.0). Please run 'update_flash' at Loader prompt to update your system firmware (1.5X3).

If...	Then...
You do not see this warning.	BIOS firmware is updated automatically if needed; go to Step 7.
You see this warning.	You must update BIOS firmware manually; go to the next step.

After the new BIOS system firmware is installed, future system firmware updates take place automatically.

Attention: The version of BIOS system firmware that ships with Data ONTAP 8.0 and higher is the minimum requirement for running this Data ONTAP release. If your system is running a lower level of system firmware, you will not be able to boot Data ONTAP 8.0 and the upgrade will fail.

6. Enter the following command:

```
update_flash
```

The system updates the firmware, displays several status messages, and displays the boot prompt.

7. At the boot prompt, enter the following command to reboot the system using the new software and, if applicable, the new firmware:

```
bye
```

8. While the HA configuration is disabled, repeat Step 4 through Step 7 at the console of system B.

Attention: Do not proceed to Step 9 until both systems in the HA configuration have been rebooted with the new version of Data ONTAP.

9. Reenable the HA configuration by entering the following command on one of the storage systems:

```
cf enable
```

Related tasks

[Installing software images from the /etc/software directory](#) on page 67

Upgrading single systems

You upgrade a single system by updating the system software and updating its firmware, then rebooting.

Before you begin

Before initiating this download procedure, verify that you have prepared for the upgrade by completing the prerequisite procedures. You must also install the Data ONTAP files to your storage system.

Note: If you are upgrading a system running Data ONTAP 7.2 or later, you can use the `software update` command to complete all or part of this procedure.

Steps

1. Choose the following option depending on whether you have already installed new system files:

If you ...	Then ...
Have already installed system files	Go to the next step.
Are installing and downloading system files in the same operation	<p>At the storage system console, enter the following command:</p> <pre>software update file_name -r</pre> <p>Then go to Step 3.</p> <p>Note: Beginning in Data ONTAP 8.0.1, the <code>-r</code> option (no automatic reboot) is the default. However, until you are running a release that supports this option, you must continue to specify the <code>-r</code> option.</p>

When you use the `software update` command without the `-d` option, the `download` command is executed by default.

2. At the system console, enter the following command to activate the new code on the storage system's boot device:

download

The `download` command provides an acknowledgment similar to the following:

```
Tue Jun 19 10:03:22 GMT [download.request:notice]:
Operator requested download initiated
download: Downloading boot device
.....
download: Downloading boot device (Service Area)
```

Then a message similar to the following appears:

```
Tues Jun 19 10:11:51 GMT [download.requestDone:notice]:
Operator requested download completed
```

Note: The storage system console is unavailable until the download procedure is complete.

Using the download procedure to activate Data ONTAP 8.0 software images takes significantly longer to complete than on earlier releases.

3. Enter the following command to shut down the storage system:

```
halt
```

After the system shuts down, the LOADER boot environment prompt appears.

4. After halting the system, check the Boot Loader messages for a warning similar to the following: Warning: The CompactFlash contains newer firmware image (1.6.0). Please run 'update_flash' at Loader prompt to update your system firmware (1.5X3).

If ...	Then ...
You do not see this warning	BIOS firmware is updated automatically if needed; go to Step 6.
You see this warning	You must update BIOS firmware manually; go to the next step.

After the new BIOS system firmware is installed, future system firmware updates take place automatically.

Attention: The version of BIOS system firmware that ships with Data ONTAP 8.0 and higher is the minimum requirement for running this Data ONTAP release. If your system is running a lower level of system firmware, you will not be able to boot Data ONTAP 8.0 and the upgrade will fail.

5. Enter the following command:

```
update_flash
```

The system updates the firmware, displays several status messages, and displays the boot prompt.

6. At the firmware environment boot prompt, enter the following command to reboot the system using the new software and, if applicable, the new firmware:

```
bye
```

Related tasks

Installing software images from the /etc/software directory on page 67

Updating firmware

Because upgrading Data ONTAP includes upgrading your firmware, you must consider the requirements for upgrading system, disk, and disk shelf firmware, as well as firmware for other components that might be installed on your system. You might also need to update firmware between Data ONTAP upgrades.

Attention: You should use the Upgrade Advisor tool (if it is available in your environment) to help you determine firmware upgrade requirements.

Next topics

[System firmware updates](#) on page 87

[Disk firmware updates](#) on page 92

[Disk shelf firmware updates](#) on page 97

[Service Processor firmware updates](#) on page 103

[RLM firmware updates](#) on page 105

[BMC firmware updates](#) on page 111

[PAM II firmware updates](#) on page 116

Related tasks

[Using the Upgrade Advisor to plan your upgrade](#) on page 19

System firmware updates

When you perform a Data ONTAP software upgrade, the firmware service image included with the Data ONTAP upgrade package is copied to your storage system's boot device. You can also update system firmware by downloading the most recent firmware for your system from the NOW site and installing the files.

If you are upgrading system firmware between Data ONTAP upgrades, you can use the nondisruptive or standard methods to update system firmware manually. You can obtain system firmware and information about how to install it from the NOW site.

Next topics

[Automatic BIOS system firmware updates](#) on page 88

[Updating system firmware nondisruptively](#) on page 88

[Updating system firmware using the standard method](#) on page 91

Related information

[System Firmware + Diagnostics Download -- now.netapp.com/NOW/cgi-bin/fw](#)

Automatic BIOS system firmware updates

Beginning with the Data ONTAP 8.0 release, the minimum BIOS release required to support Data ONTAP also enables automatic BIOS updates.

After the minimum version is running, subsequent updates take place automatically during the boot sequence whenever Data ONTAP detects that a version resident on the boot device is more recent than the running version.

However, to update firmware from an earlier version to the latest version available, you must run the `update_flash` command manually from the boot prompt on the system being upgraded.

Subsequent system firmware updates are automatic.

Attention: Your system must be running the minimum required version or later to complete the upgrade to Data ONTAP 8.0 or later. If required firmware is not resident on the boot device, Data ONTAP 8.0.x releases will not boot and the upgrade will fail.

The following are the minimum BIOS system firmware versions required to support Data ONTAP.

Platform	Minimum version
60xx	BIOS 1.5X2 or later
31xx	
3070	BIOS 2.2X1 or later
3040	
2040	BIOS/NABL 6.0 or later

62xx and 32xx platforms ship with the minimum system firmware versions. All subsequent firmware updates are automatic. It is not necessary to run the `update_flash` command on these platforms for normal system firmware updates.

Updating system firmware nondisruptively

The nondisruptive update method is appropriate when you need to maintain service availability during the firmware update.

Before you begin

You should ensure that your HA configuration is functioning correctly and meets the requirements for nondisruptive upgrades.

You must download firmware from the NOW site on your Windows or UNIX client or your HTTP server before you begin this procedure.

Steps

1. Obtain the firmware download files using the `software update` command, following directions on the NOW site.
2. On each storage system, referred to as system A and system B in the following steps, enter the following command as directed:

```
priv set advanced
```

The asterisk (*) after the storage system name indicates 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. Take one of the following actions:

If CIFS...	Then...
Is not in use in system A.	Go to Step 5.
Is in use in system A.	<ol style="list-style-type: none"> a. Enter the following command: <pre>cifs terminate -t nn</pre> <i>nn</i> is a notification (in seconds) appropriate for your clients. After that period of time, proceed to Step 5. b. Wait for <i>nn</i> seconds and then go to Step 5.

5. If the automatic giveback option (`cf.giveback.auto.enable`) is set to `on`, disable automatic giveback by entering the following command on one of your storage systems in the high-availability configuration:

```
options cf.giveback.auto.enable off
```

After the upgrade procedure, reset this option to `on` (if desired).

6. At the console of system B, enter the following command:

```
cf takeover
```

This command causes system A to shut down gracefully and leaves system B in takeover mode.

7. To display the LOADER boot prompt at the system A console, press Ctrl-C at the system A console when instructed after the boot sequence starts.

You can also display the LOADER prompt by pressing Ctrl-C at the system A console when the "Waiting for giveback" message appears at the console of system A. When prompted to halt the node rather than wait, enter `y`.

8. After halting the node, check the Boot Loader messages for a warning similar to the following:
Warning: The CompactFlash contains newer firmware image (1.6.0). Please run 'update_flash' at Loader prompt to update your system firmware (1.5X3).

If you...	Then ...
Do not see this warning.	BIOS firmware is updated automatically if needed; go to Step 12.
See this warning.	You must update BIOS firmware manually; go to the next step.

After the new BIOS system firmware is installed, future system firmware updates take place automatically.

Attention: The version of BIOS system firmware that ships with Data ONTAP 8.0 and higher is the minimum requirement for running this Data ONTAP release. If your system is running a lower level of system firmware, you will not be able to boot Data ONTAP 8.0 and the upgrade will fail.

- At the boot prompt, enter the following command to reset the system:

```
bye
```

- Display the LOADER boot prompt again at the system A console by repeating Step 7.

- Enter the following command:

```
update_flash
```

The system updates the firmware, displays several status messages, and displays the boot prompt.

- Enter the following command to reboot the system using the new firmware and software:

```
bye
```

- Choose the option that describes your configuration.

If FCP or iSCSI...	Then when the "Waiting for giveback" message appears on the console of system A...
Is not in use in system A	Enter the following command at the console of system B: cf giveback
Is in use in system A	Wait for at least eight minutes to allow host multipathing software to stabilize, then enter the following command at the console of system B: cf giveback

Attention: The `cf giveback` command can fail because of open client sessions (such as CIFS sessions), long-running operations, or operations that cannot be restarted (such as tape backup or SyncMirror resynchronization). If the `cf giveback` command fails, terminate any CIFS session or long-running operations gracefully (because the `-f` option will immediately terminate any CIFS sessions or long-running operations) and then enter the following command (with the `-f` option):

```
cf giveback -f
```

For more information about the behavior of the `-f` option, see the `cf(1)` man page.

The command causes system A to reboot with the new firmware and resume normal operation as a high-availability partner.

14. Choose the following option that describes your configuration.

If FCP or iSCSI...	Then...
Is not in use in system A	Repeat Step 4 through Step 14 to update the partner storage system; that is, bring down and update system B with partner A in takeover mode.
Is in use in system A	<p>a. After system A resumes normal operation as a high-availability partner, wait for at least eight minutes to allow host multipathing software to stabilize.</p> <p>b. Repeat Step 4 through Step 12 to update the partner storage system; that is, bring down and update system B with system A in takeover mode.</p>

After you finish

If desired, reenable automatic giveback.

Updating system firmware using the standard method

The standard firmware update method is appropriate when you can schedule downtime for the system firmware update.

Before you begin

You must obtain the system firmware from the NOW site on your Windows or UNIX client or your HTTP server before you begin this procedure.

Steps

1. On each system you are upgrading, enter the following command:

```
priv set advanced
```

The asterisk (*) after the storage system name indicates that you are in advanced mode.

2. 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.

3. On either system, disable the HA configuration by entering the following command:

```
cf disable
```

4. Continue installing the firmware on each system by following directions from the NOW site.

5. Reenable the HA configuration by entering the following command on one of the systems:

```
cf enable
```

Disk firmware updates

You should update to the latest disk firmware version when you upgrade Data ONTAP. In some upgrade scenarios, disk firmware updates are mandatory.

Next topics

How disk firmware is updated on page 92

Service availability during disk firmware updates on page 93

Detecting outdated disk firmware on page 95

When to update disk firmware manually on page 96

Command for updating disk firmware on page 96

How disk firmware is updated

Disk firmware is automatically updated in certain circumstances.

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

- You add new disks or a disk shelf.

Note: When hot-adding SAS shelves, firmware is not updated automatically. You must manually check and update any out-of-date drive, shelf, and ACP firmware.

- Data ONTAP detects disk firmware updates in the `/etc/disk_fw` directory. Data ONTAP scans the `/etc/disk_fw` directory for new disk firmware every two minutes.

Each storage system is shipped with an `/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 update disk firmware.
- During a manual firmware update
 - You might be directed to download a disk firmware update 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 updates 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.

Service availability during disk firmware updates

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

Disk firmware updates can take place in two ways:

- Background (nondisruptive) disk firmware update

Nondisruptive disk firmware updates 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)
- Standard disk firmware update

In Data ONTAP 7.2 and later, disk firmware updates for RAID4 aggregates must complete before the new Data ONTAP version can finish booting. Storage system services are not available until the disk firmware update finishes.

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

Next topics

[Verifying RAID protection type](#) on page 93

[Understanding background disk firmware updates](#) on page 94

[Understanding standard disk firmware updates](#) on page 95

Verifying RAID protection type

You should check the RAID type of your root volume before you update its firmware, because if any volume, including the root volume, is configured with RAID4 protection, a standard disk firmware update (interrupting storage system services) will take place at the next reboot when new disk firmware is present on the system.

Step

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

```
aggr status
```

You see output similar to the following:

Aggr	State	Status	Options
data2_vol	online	raid-dp,	flex

data1_vol	online	raid-dp, flex	
vol0	online	raid4, flex	root

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 update its firmware, and reconfigure it if necessary, if you require a nondisruptive disk firmware update.

Understanding background disk firmware updates

There are many important issues to consider when performing a background disk firmware update.

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 update 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 update, the storage system functions normally. You will see status messages as disks are taken offline to update firmware and brought back online when the firmware update is complete. Background disk firmware updates proceed sequentially for active data disks and for spare disks. Sequential disk firmware updates 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 reconstruction operations. However, a disk would normally remain offline for a very short time (a few minutes at most) and therefore would not interfere with normal system operation.

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

- Degraded volumes are 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 updates will resume when these conditions are addressed. For more information about determining volume status and state, see the *Data ONTAP 7-Mode Storage Management Guide*.

Automatic background disk firmware updates are overridden when the `disk_fw_update` command is issued.

Note: Automatic background disk firmware updates 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 high-availability takeover, when the `disk_fw_update` command is issued, or when a disk firmware update is required for disks in a RAID4 aggregate. You are advised not to change the default value unless you are directed to by technical support.

Related concepts

[Command for updating disk firmware](#) on page 96

Understanding standard disk firmware updates

During a standard disk firmware update, the disks of the affected drive types are not available.

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 updates 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 updates can take many minutes depending on the number of drive types and disk drives per type on your storage system.

Note: If you upgrade RAID protection to RAID-DP, disk firmware updates take place in the background and are nondisruptive.

Detecting outdated disk firmware

AutoSupport messages include information about disk firmware installed on your storage system. The NOW Installed Systems 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` directory has been superseded.

Before you begin

To use the NOW Installed Systems service to monitor disk firmware versions, your storage system must meet the following requirements:

- AutoSupport must be enabled on your storage system.
For more information about AutoSupport, see the *Data ONTAP 7-Mode System Administration Guide*.
- You must have registered your NetApp products.

About this task

These AutoSupport notices indicate that the disk firmware on at least some of your disks will be updated during your next Data ONTAP upgrade, which can help you plan your upgrade.

Steps

1. Use a Web browser to go to the NOW site at now.netapp.com.
2. Select **My Support > View Installed Systems**.
3. Display the product details for the storage system 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 **Health Check Details**.

Result

If there is a firmware update available for your storage system, you will see a message with a link to a Firmware Analysis page. If there is a notice on the Firmware Analysis page that newer disk firmware is available for your system, a disk firmware update 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.

After you finish

Determine if you should update your disk firmware now.

Related concepts

When to update disk firmware manually on page 96

Related information

My Support: Systems - now.netapp.com/eservice/Systems.jsp

When to update disk firmware manually

If you receive error messages about firmware compatibility, you must manually update your disk firmware.

After downloading new disk firmware from the NOW site, you must enter the `disk_fw_update` command at the storage system prompt to install the new firmware.

You must also update disk firmware manually if you hot-add SAS shelves.

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.

Related concepts

Command for updating disk firmware on page 96

Command for updating disk firmware

You need to use the `disk_fw_update` command from the storage system console to update firmware on all disks or a on specified disk on a storage system.

The `disk_fw_update` command updates disks for which firmware files are present in the `/etc/disk_fw` directory and which need to be updated. It does not update other disks.

The `disk_fw_update` command is applicable to SCSI, Fibre Channel, SATA, and SAS disks.

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

Attention: 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 update that is in progress.

Disk shelf firmware updates

You should update to the latest disk shelf firmware version when you upgrade Data ONTAP. In some upgrade scenarios, disk shelf firmware updates are mandatory.

Note: Disk shelf firmware updates are mandatory when hot-adding a disk shelf. See your disk shelf documentation for more information.

Next topics

[How disk shelf firmware is updated](#) on page 97

[Service availability during disk shelf firmware updates](#) on page 98

[Detecting outdated disk shelf firmware](#) on page 99

[Updating disk shelf firmware manually](#) on page 100

[Updating ACP firmware](#) on page 102

How disk shelf firmware is updated

When you upgrade Data ONTAP, disk shelf firmware (firmware for modules on disk shelves) is updated automatically if the firmware on the shelves is older than the firmware that is bundled with the Data ONTAP system files. You can also update disk shelf firmware by downloading the most recent firmware for your shelf modules from the NOW site and installing the files.

The module (AT series, ESH series, or SAS) in a disk shelf provides for the interconnect of the disks to the host bus adapter interface, including signal integrity when disks are swapped. In AT- and ESH-based shelves, there are two modules in the middle of the rear of the disk shelf, one for Channel A and one for Channel B. SAS modules are internal components in FAS2040 systems. Updated firmware for these modules is made available periodically.

Each storage system is shipped with an `/etc/shelf_fw` directory that contains the latest disk shelf firmware versions available at that time.

Disk shelf firmware updates can be added to this directory at the following times:

- After a Data ONTAP upgrade
Disk shelf firmware updates are often included in Data ONTAP upgrade packages. If the version in `/etc/shelf_fw` is higher than the installed version, the new version will be downloaded and installed during the `reboot` or `cf giveback` phase as part of the Data ONTAP upgrade process.
- During a manual firmware update
You might need to download a disk shelf firmware update from the NOW site if you plan to perform a nondisruptive upgrade of Data ONTAP software, or if you receive a notice from NetApp.
- When you hot-add a SAS shelf

Data ONTAP scans the `/etc/shelf_fw` directory for new firmware every two minutes (on systems with software-based disk ownership). If new disk shelf firmware is detected—that is, if there is a disk

shelf firmware file in the `/etc/shelf_fw` directory that has a higher revision number than the current firmware on the shelf module—the new firmware is automatically downloaded to the disk shelf module.

The following events in Data ONTAP can also trigger an automatic disk shelf firmware update when there is new firmware in the `/etc/shelf_fw` directory:

- The `reboot` command is issued.
- The `cf giveback` command is issued.
- New disk drives are inserted.
- New shelf modules are inserted.
- NetApp Health Trigger (NHT) AutoSupport messages are sent.

For more information about disk shelves and disk shelf modules, see the *Data ONTAP 7-Mode High-Availability Configuration Guide* and the *Hardware and Service Guide* for your shelves.

Service availability during disk shelf firmware updates

When you upgrade to the current Data ONTAP release, the availability of storage system services during a disk shelf firmware update depends on the type of shelf modules your system uses.

The following table summarizes Data ONTAP service availability during disk shelf firmware updates for these modules:

Module	Disk shelf model	System downtime required?
AT-FCX	DS14mk2 AT	With Multipath Storage and firmware version 37: No Note: Multipath Storage can be implemented in HA pairs or standalone systems. AT-FCX firmware can be upgraded nondisruptively in either configuration. Without Multipath Storage: Yes
AT-FC2	DS14mk2 AT	Yes
AT-FC		
ESH4	DS14mk4 FC or DS14mk2 FC	No
ESH2	DS14mk2 FC	

Module	Disk shelf model	System downtime required?
SAS	DS4243	No
	DS2246	No
	FAS20xx or SA200 internal shelves	With firmware version 0500 and later: No With firmware version 0400 and earlier: Yes

Attention:

You cannot use the nondisruptive method to upgrade Data ONTAP under the following circumstances:

- AT-FCX disk shelves are attached to your system, unless you use Multipath Storage and unless the firmware for these modules is version 37 or higher.
- You have AT-FC, or AT-FC2-based disk shelves attached to your system.
- You have internal SAS modules in a FAS20xx or SA200 system, unless the firmware for these modules is version 0500 or higher.

Detecting outdated disk shelf firmware

If you want to perform a nondisruptive upgrade of Data ONTAP software when there are AT-based disk shelves attached to your system, or if you are directed to update disk shelf firmware, you must find out what firmware is installed on disk shelves attached to your system.

Before you begin

You should use the Upgrade Advisor (if it is available in your environment) to assess the status of your disk shelf firmware before upgrading.

Steps

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

```
sysconfig -v
```

2. Locate the shelf information in the `sysconfig -v` output.

Example

```
Shelf 1: AT-FCX Firmware rev. AT-FCX A: 36 AT-FCX B: 36
Shelf 2: AT-FCX Firmware rev. AT-FCX A: 36 AT-FCX B: 36
```

3. Go to the disk shelf firmware information on the NOW site and determine the most recent firmware version for your shelves.

4. Take the appropriate action.

If the disk shelf firmware version in the <code>sysconfig -v</code> output is ...	Then ...
The same as the most recent version on the NOW site	No disk shelf firmware update is required at this time.
Earlier than the most recent version on the NOW site	Update your disk shelf firmware manually.

Related tasks

Using the Upgrade Advisor to plan your upgrade on page 19

Related information

Disk Shelf Firmware on NOW -- now.netapp.com/NOW/download/tools/diskshelf/

Updating disk shelf firmware manually

You must run the `storage download shelf` command after downloading new disk shelf firmware from the NOW site.

About this task

By running the `storage download shelf` command once, you upgrade all eligible modules connected to both controllers in HA configurations.

The command updates the modules sequentially:

- ESH series and SAS IOM series
The command begins with the module that is currently reporting SCSI Enclosure Services (SES) status.
- AT series and SAS modules in FAS2040 systems
The command first updates all A modules, then all B modules.

Attention: Do not place firmware files in the `/etc/shelf_fw` directory unless you intend to update disk shelf firmware immediately. Several events in Data ONTAP can trigger an automatic disk shelf firmware update if there is a disk shelf firmware file in the `/etc/shelf_fw` directory that has a higher revision number than the current firmware on the shelf module.

Do not use the nondisruptive method (that is, the `cf takeover` and `cf giveback` commands) to update disk shelf firmware. Doing so will prevent access to data on disk shelves for a much longer period than using the `storage download shelf` command.

Steps

1. Find and download the most recent firmware for your shelves on the NOW site.

- Follow the instructions on the NOW site to extract your firmware files to the `/etc/shelf_fw` directory in the root volume of your storage system.
- Choose the following option that describes your configuration.

If you are running CIFS on systems with one of the following configurations ...	Then ...
<ul style="list-style-type: none"> ESH-based disk shelves DS2246 disk shelves DS4243 disk shelves AT-FCX-based disk shelves running firmware version 37 or higher FAS2040 internal shelves 	Go to the next step.
<ul style="list-style-type: none"> AT-based disk shelves AT-FCX-based disk shelves running firmware version 36 or lower 	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 the next step.

- Enter the following command at the storage system console to access the advanced administrative commands:

```
priv set advanced
```

The prompt now displays an asterisk (*) after the storage system name to indicate that you are in the advanced mode.

- Depending on your upgrade scenario, enter one of the following commands to upgrade the disk shelf firmware.

If you want to upgrade the disk shelf firmware on ...	Then enter the following command at the storage system console:
All the disk shelves in your system	storage download shelf
The shelves attached to a specific adapter	storage download shelf adapter_name

- To confirm that you want to upgrade the firmware, enter the following key:

```
y
```

- Enter the following command to verify the new disk shelf firmware:

```
sysconfig -v
```

- Enter the following command to return to the standard administrative console prompt:

```
priv set admin
```

9. If you terminated CIFS before updating shelf firmware, reenable it by entering the following command:

```
cifs restart
```

Related information

Disk Shelf Firmware on NOW -- now.netapp.com/NOW/download/tools/diskshelf/

Updating ACP firmware

If your disk shelves include Shelf Alternate Control Path Management (ACP) functionality, you can update ACP firmware by running the `storage download acp` command after downloading new ACP processor firmware from the NOW site.

Before you begin

ACP interfaces must be cabled properly and ACP software must be configured correctly. For more information, see the *Data ONTAP 7-Mode Storage Management Guide* and the *Installation and Service Guide* for your disk shelf.

About this task

When you upgrade Data ONTAP, ACP firmware (firmware for ACP processors on disk shelves) is updated automatically if the firmware in the ACP processors is older than the firmware that is bundled with the Data ONTAP system files. However, it might be necessary to update ACP firmware (for example, when hot-adding a disk shelf) by downloading the most recent firmware from the NOW site and installing the files.

Note: Installing ACP firmware can take several minutes, but it will not disrupt client access during that time. However, normal ACP recovery capabilities will not be available while the firmware upgrade is in progress.

Steps

1. Find and download the most recent ACP firmware on the NOW site.
2. Follow the instructions on the NOW site to extract your firmware files to the `/etc/acpp_fw` directory in the root volume of your storage system.
3. Enter the following command to update the ACP firmware:

```
storage download acp
```

For more information about the command, see the `storage(1)` man page.

4. Enter the following command to verify the new ACP firmware:

```
storage show acp
```

You should see command output similar to the following while the ACP firmware is being updated:

```

Alternate Control Path: Enabled
Ethernet Interface:     e0c
ACP Status:            Active
ACP IP Address:        192.168.0.67
ACP Domain:            192.168.0.0
ACP Netmask:           255.255.252.0
ACP Connectivity Status: Full Connectivity

```

Shelf_Module	Reset_Cnt	IP_Address	FW_Version	Module_Type	Status
8a.00.A	000	192.168.2.60	01.10	IOM6	inactive (upgrading firmware)
8a.00.B	000	192.168.2.112	02.00	IOM6	active
8a.02.A	000	192.168.1.218	01.10	IOM3	active
8a.02.B	000	192.168.1.78	01.10	IOM3	active
8a.10.A	000	192.168.3.77	01.10	IOM3	active
8a.10.B	000	192.168.3.83	01.10	IOM3	active

When the update has completed, you will see output similar to the following when you reissue the command:

Shelf_Module	Reset_Cnt	IP_Address	FW_Version	Module_Type	Status
8a.00.A	000	192.168.2.60	02.00	IOM6	active
8a.00.B	000	192.168.2.112	02.00	IOM6	active
8a.02.A	000	192.168.1.218	01.10	IOM3	active
8a.02.B	000	192.168.1.78	01.10	IOM3	active
8a.10.A	000	192.168.3.77	01.10	IOM3	active
8a.10.B	000	192.168.3.83	01.10	IOM3	active

Related information

[Disk Shelf Firmware on NOW -- now.netapp.com/NOW/download/tools/diskshelf/](http://now.netapp.com/NOW/download/tools/diskshelf/)

Service Processor firmware updates

Service Processor (SP) is a remote management device that is included in 32xx and 62xx systems. You can upgrade the SP firmware by downloading and updating the SP firmware using the Data ONTAP CLI or the SP CLI.

For information about what the SP is and how it works, see the *Data ONTAP 7-Mode System Administration Guide*.

Next topics

[Using the Data ONTAP CLI to update the SP firmware](#) on page 104

[Using the SP CLI to update the SP firmware](#) on page 104

Using the Data ONTAP CLI to update the SP firmware

You can update the SP firmware at the storage system prompt.

Before you begin

You must have the following items before you can download and update the firmware:

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

Steps

1. Go to Firmware Instructions for the Service Processor at the NOW site.
2. Click the `SP_FW.zip` link to download the file from the NOW site to your HTTP server.
3. At the storage system prompt, enter the following command:

```
software update http://Web_server/SP_FW.zip -f
```
4. When the `software update` command is finished, enter the following command at the storage system prompt:

```
sp update
```
5. When the system prompts you to update SP, enter `y` to continue.

Result

SP is updated and you are prompted to reboot SP. Wait approximately 60 seconds to allow SP to reboot.

Note: If your console connection is not through SP, the connection remains active during the SP reboot.

If your console connection is through SP, you lose your console connection to the storage system. In approximately one minute, SP reboots and automatically re-establishes the connection.

Related information

[Service Processor \(SP\) Firmware -- now.netapp.com/NOW/download/tools/sp_fw/](http://now.netapp.com/NOW/download/tools/sp_fw/)

Using the SP CLI to update the SP firmware

You can update the SP firmware at the SP prompt.

Before you begin

You must have the following items before you can download and update the firmware:

- Access to a Web server on a network accessible to your storage system
- The name and IP address of the Web server
- Access to the storage system SP CLI

Steps

1. Go to Firmware Instructions for the Service Processor at the NOW site.
2. Click the `SP_FM.tar.gz` link to download the file from the NOW site to your HTTP server.
3. Log in to SP by entering the following command at the administration host:

```
ssh username@SP_IP_address
```

4. At the SP prompt, enter the following command:

```
sp update http://Web_server_addr/SP_FW.tar.gz
```

5. When you are prompted to reboot SP, enter the following command at the SP prompt:

```
sp reboot
```

Related information

[Service Processor \(SP\) Firmware -- now.netapp.com/NOW/download/tools/sp_fw/](http://now.netapp.com/NOW/download/tools/sp_fw/)

RLM firmware updates

You can upgrade the Remote LAN Module (RLM) firmware by downloading and updating the RLM firmware using the Data ONTAP CLI or the RLM CLI.

For information about what the RLM is and how it works, see the *Data ONTAP 7-Mode System Administration Guide*.

Next topics

[Requirements for RLM firmware version 4.0 and later](#) on page 105

[Using the Data ONTAP CLI to update the RLM firmware](#) on page 106

[Using the RLM CLI to update the RLM firmware](#) on page 108

[RLM firmware update problems](#) on page 110

Requirements for RLM firmware version 4.0 and later

RLM firmware versions 4.0 and later require a different layout on flash media. You must ensure that you are running the latest 3.1.x RLM firmware to enable the transition to the new layout, then update to the 4.0 or later firmware.

You must be running the latest 3.1.x to update to 4.0. If you are running a firmware version earlier than 3.1, you must first perform an intermediate update to the latest 3.1.x firmware, then update from 3.1 to 4.0 in a separate operation.

Attention: Regardless of whether you update RLM firmware from the Data ONTAP CLI or the RLM CLI, *do not* update directly from a firmware version earlier than 3.1 to 4.0 or later. Doing so will corrupt the RLM flash device.

If you are updating to version 4.0 or later from either the Data ONTAP CLI or the RLM CLI, you must run the `rlm update` command with the `-f` option for a full image update. Further updates do not require the `-f` option.

If you are updating RLM firmware from the RLM CLI, you can use the normal procedure.

For information about configuring the RLM, see the *Data ONTAP 7-Mode System Administration Guide*.

Using the Data ONTAP CLI to update the RLM firmware

You can update RLM firmware at the storage system prompt.

Before you begin

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

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

Steps

1. Enter the following command to display the current RLM firmware version:

```
rlm status
```

You see a display similar to the following:

```
Remote LAN Module           Status: Online
  Part Number:              000-00000
  Revision:                 A0
  Serial Number:            00000
  Firmware Version:         1.2
  Mgmt MAC Address:         00:00:00:00:00:00
  Ethernet Link:            up
  Using DHCP:               no
```

2. Complete the steps as directed in the following table based on your RLM firmware version.

If the firmware version is...	Then...
Earlier than 3.1	Complete Steps 3 through 7 to upgrade to the latest 3.1.x version. If you want to update to version 4.0 or later, you must also complete Steps 8 through 13.
3.1.x	Complete Steps 8 through 13.

If the firmware version is... Then...

4.0 or later Complete Steps 3 through 7.

3. Go to Firmware Instructions for the Remote LAN Module at the NOW site.
4. Click the `RLM_FM.zip` link to download the file from the NOW site to your HTTP server.
You should download the latest 3.1.x or 4.0 firmware, depending on the update that is required.
If the latest 4.x firmware on the NOW site is the same as the version running on your RLM, it is not necessary to update RLM firmware at this time.
5. Enter the following command at the storage system prompt:
software update http://Web_server/RLM_FW.zip -f
6. When the `software update` command is finished, enter the following command at the storage system prompt:
rlm update
Messages inform you of the progress of the update.
7. When the system prompts you to update RLM, enter `y` to continue.
RLM is updated and you are prompted to reboot RLM. Wait approximately 60 seconds to allow RLM to reboot.

Note: If your console connection is not through RLM, it stays active during reboot.

If...	Then...
You have already updated to firmware version 4.0, or you are not planning to update to 4.0.	The procedure is complete.
You are updating firmware to version 4.0 or higher for the first time.	Proceed to the next step.

8. If you have not already done so, download the version 4.0 firmware as described in Steps 3 and 4.
9. Enter the following command at the storage system prompt:
software update http://Web_server/RLM_FW.zip -f
10. When the `software update` command is finished, enter the following command at the storage system console to access the advanced administrative commands:
priv set advanced
The prompt now displays an asterisk (*) after the storage system name to indicate that you are in the advanced mode.
11. Enter the following command at the storage system prompt:
rlm update -f

Note: Be sure to use the `-f` option to enable the new flash layout for IPv6.

Messages inform you of the progress of the update.

- When the system prompts you to update RLM, enter `y` to continue.

RLM is updated and you are prompted to reboot RLM. Wait approximately 60 seconds to allow RLM to reboot.

Note: If your console connection is not through RLM, it stays active during reboot.

- Enter the following command to return to the standard administrative console prompt:

```
priv set admin
```

Related information

Remote LAN Module (RLM) Firmware -- now.netapp.com/NOW/download/tools/rlm_fw/

Using the RLM CLI to update the RLM firmware

You can update the RLM firmware at the RLM prompt.

Before you begin

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

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

Steps

- Enter the following command to display the current RLM firmware version:

```
rlm status
```

You see a display similar to the following:

```
Remote LAN Module           Status: Online
  Part Number:              000-00000
  Revision:                 A0
  Serial Number:            00000
  Firmware Version:         1.2
  Mgmt MAC Address:         00:00:00:00:00:00
  Ethernet Link:            up
  Using DHCP:               no
```

- Complete the steps as directed in the following table based on your RLM firmware version.

If the firmware version is...	Then...
Earlier than 3.1	Complete Steps 3 through 7 to upgrade to the latest 3.1.x version. If you want to update to version 4.0 or later, you must also complete Steps 8 through 13.
3.1.x	Complete Steps 8 through 11.
4.0 or later	Complete Steps 3 through 7.

- Go to Firmware Instructions for the Remote LAN Module at the NOW site.
- Click the `RLM_FM.tar.gz` link to download the file from the NOW site to your HTTP server.
You should download the latest 3.1.x or 4.0 firmware, depending on the update that is required.
If the latest 4.x firmware on the NOW site is the same as the version running on your RLM, it is not necessary to update firmware at this time.

- Log in to the RLM by entering the following command at the administration host:

```
ssh username@RLM_IP_address
```

- Enter the following command at the RLM prompt:

```
rlm update http://Web_server_addr/RLM_FW.tar.gz
```

- When you are prompted to reboot the RLM, enter the following command at the RLM prompt:

```
rlm reboot
```

Note: If your console connection is through the RLM, you lose your console connection to the storage system. In approximately one minute, the RLM reboots and automatically re-establishes the connection.

If...	Then...
You have already updated to firmware version 4.0, or you are not planning to update to 4.0.	The procedure is complete.
You are updating firmware to version 4.0 or higher for the first time.	Proceed to the next step.

- If you have not already done so, download the version 4.0 firmware as described in Steps 3 and 4.

- Enter the following command at the RLM prompt:

```
rlm update -f http://Web_server_addr/RLM_FW.tar.gz
```

- When you are prompted to reboot the RLM, enter the following command at the RLM prompt:

```
rlm reboot
```

Note: If your console connection is through the RLM, you lose your console connection to the storage system. In approximately one minute, the RLM reboots and automatically re-establishes the connection.

Related information

[Remote LAN Module \(RLM\) Firmware -- now.netapp.com/NOW/download/tools/rlm_fw/](http://now.netapp.com/NOW/download/tools/rlm_fw/)

RLM firmware update problems

A RLM firmware update failure can occur for a number of reasons. You can troubleshoot a firmware failure by searching for EMS events.

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

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

You can troubleshoot a firmware failure by searching for EMS events. For instance, an error message indicating that the firmware update failed due to a communication error might appear as follows:

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

For more information about the Event Management System (EMS), see the `ems(1)` man page.

Next topics

[Troubleshooting RLM firmware update problems with the Data ONTAP CLI](#) on page 110

[Troubleshooting RLM firmware update problems with the RLM CLI](#) on page 111

Troubleshooting RLM firmware update problems with the Data ONTAP CLI

You can troubleshoot a firmware update using the Data ONTAP CLI.

Steps

1. Verify that RLM is online by entering the following command at the storage system prompt:


```
rlm status
```
2. Update the RLM firmware by following the instructions described in "Using the Data ONTAP CLI to update the RLM firmware."
3. Verify that you are using the correct filename (`filename.zip`) of the RLM firmware.
4. Reboot RLM by entering the following command at the storage system prompt:


```
rlm reboot
```

It takes approximately one minute for the RLM to reboot.

5. If the RLM does not reboot after one minute, repeat Steps 1 through 4.

If the RLM still does not reboot, contact technical support for assistance.

Related tasks

[Using the Data ONTAP CLI to update the RLM firmware](#) on page 106

Troubleshooting RLM firmware update problems with the RLM CLI

You can troubleshoot a firmware update using the RLM CLI.

Steps

1. Verify that RLM is online by entering the following command at the storage system prompt:

```
rlm status
```

2. From a browser, access the RLM firmware file on your Web server.
3. Verify that you are using the correct filename (*filename.tar.gz*) of the RLM firmware.
4. Update the firmware by entering the following command at the RLM prompt:

```
rlm update http://path_hostname/RLM.FW.tar.gz [-f]
```

If this command fails, replace *path_hostname* with the corresponding IP address.

The `-f` option issues a full image update.

5. Reboot RLM by entering the following command at the storage system prompt:

```
rlm reboot
```

Related tasks

[Using the RLM CLI to update the RLM firmware](#) on page 108

BMC firmware updates

Baseboard Management Controller (BMC) firmware is bundled with the Data ONTAP software image. When you perform a Data ONTAP software upgrade on a system with a BMC, the BMC firmware included with the Data ONTAP upgrade image is installed on your storage system's boot device if the firmware in the image is a later version than the firmware on your system.

If new BMC firmware was installed, you must run the `update_bmc` boot-loader macro to load the new BMC firmware on the BMC device. You can load the BMC firmware using the nondisruptive method in HA configurations, or you can use the standard method in both high-availability and single-system configurations.

For information about what the BMC is and how it works, see the *Data ONTAP 7-Mode System Administration Guide*.

Next topics

[Detecting outdated BMC firmware](#) on page 112

[Updating BMC firmware nondisruptively](#) on page 113

[Updating BMC firmware using the standard method](#) on page 115

Related concepts

[Installing Data ONTAP software images](#) on page 63

Detecting outdated BMC firmware

After upgrading Data ONTAP software, you should determine if new BMC firmware was loaded onto your system.

Steps

1. At the storage system prompt, enter the following command to identify the currently installed BMC firmware version:

```
bmc status
```

Example

```
storage_system> bmc status
      Baseboard Management Controller:
      Firmware Version: 1.1
```

2. At the storage system prompt, enter the following command to identify the version of the BMC firmware on the boot device:

```
version -b
```

The console displays the contents of the boot device's File Allocation Table (FAT) file system, including the BMC firmware version.

Example

```
storage_system> version -b
1:/x86_elf/kernel/primary.krn: OS 7.2.2L1X9
1:/backup/x86_elf/kernel/primary.krn: OS Rgb-shuarN_070510_0030
1:/x86_elf/diag/diag.krn: 4.8
1:/x86_elf/firmware/deux/firmware.img: Firmware 3.1.0
1:/x86_elf/firmware/SB_XIV/firmware.img: BIOS/NABL Firmware 3.0
1:/x86_elf/firmware/SB_XIV/bmc.img: BMC Firmware 1.0
```

3. Compare the output of the `bmc status` and `version -b` commands.

If ...	Then ...
The commands show the same BMC firmware version	No BMC firmware update is required at this time.
The BMC firmware version in the <code>version -b</code> output is later than the version in the <code>bmc status</code> status	Use the nondisruptive or standard method to update BMC firmware.

Updating BMC firmware nondisruptively

The nondisruptive update method is appropriate when you need to maintain service availability during BMC firmware updates. To use this method, your storage systems must be in HA configurations.

Before you begin

You must have determined if new BMC firmware is present on your system before performing this procedure.

Steps

1. On each storage system, referred to as system A and system B in the following steps, enter the following command:

```
priv set advanced
```

The prompt displays an asterisk (*) after the storage system name to indicate that you are in advanced mode.

2. Take one of the following actions:

If CIFS...	Then...
Is not in use in system A.	Go to Step 3.
Is in use in system A.	Enter the following command: cifs terminate -t nn <i>nn</i> is a notification period (in minutes) appropriate for your clients after which CIFS services are terminated. After that period of time, proceed to Step 3.

3. If the automatic giveback option (`cf.giveback.auto.enable`) is set to `on`, disable automatic giveback by entering the following command on one of your systems in the HA configuration:

```
options cf.giveback.auto.enable off
```

After the upgrade procedure, reset this option to `on` (if desired).

4. At the console of system B, enter the following command:

```
cf takeover
```

This command causes system A to shut down gracefully and leaves system B in takeover mode.

5. To display the `LOADER` boot prompt at the system A console, press `Ctrl-C` at the system A console when instructed after the boot sequence starts.

You can also display the `LOADER` prompt by pressing `Ctrl-C` at the system A console when the "Waiting for giveback" message appears at the console of system A. When prompted to halt the node rather than wait, enter `y`.

- At the `LOADER` prompt, enter the following command to reset the system:

```
bye
```

- Display the `LOADER` boot prompt again at the system A console by repeating Step 5.
- Enter the following command from the `LOADER` prompt:

```
update_bmc
```

The `update_bmc` macro updates the BMC firmware from the image on the boot device and displays a message on the console.

Example

```
LOADER> update_bmc
BMC firmware version: 1.2
Programming: this might take up to 120 seconds to complete...

pre-init time          [bmc.reset.power:notice]: Hard reset by
external power-cycle.
BMC Release 1.2
Press ^G to enter BMC command shell

Important: In order for the BMC firmware changes to fully take effect,
it is necessary to reboot using the "bye" command before starting ONTAP
```

If the new BMC firmware also has a new non-volatile memory management (NVMEM) battery firmware image, the battery firmware is updated automatically.

- Enter the following command to reboot the storage system using the new firmware and software:

```
bye
```

- When the "Waiting for giveback" message appears on the console of system B, enter the following command:

```
cf giveback
```

This command causes system A to reboot with the new firmware and resume normal operation as the HA configuration partner.

- Repeat Step 2 through Step 10 to update the partner system; that is, bring down and update system B with partner A in takeover mode.

- Enter the following command to return to the standard administrative console prompt:

```
priv set admin
```

Updating BMC firmware using the standard method

The standard firmware update method is appropriate when you can schedule downtime for system firmware updates.

Before you begin

You must have determined if new BMC firmware is present on your system before performing this procedure.

Steps

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

```
halt
```

The storage system console displays the boot environment prompt.

2. Enter the following command from the LOADER prompt:

```
update_bmc
```

The `update_bmc` macro updates the BMC firmware from the image on the boot device and displays a message on the console.

Example

```
LOADER> update_bmc
BMC firmware version: 1.2
Programming: this might take up to 120 seconds to complete...
```

```
pre-init time          [bmc.reset.power:notice]: Hard reset by
external power-cycle.
BMC Release 1.2
Press ^G to enter BMC command shell
```

Important: In order for the BMC firmware changes to fully take effect, it is necessary to reboot using the "bye" command before starting ONTAP

If the new BMC firmware also has a new non-volatile memory management (NVMEM) battery firmware image, the battery firmware is updated automatically.

3. After the BMC firmware is updated, enter the following command from the LOADER prompt to restart the system:

```
bye
```

PAM II firmware updates

Firmware for Performance Acceleration Module II (PAM II) devices is included with distribution files for Data ONTAP upgrades. If the running firmware is older than the firmware that is bundled with the Data ONTAP system files, it is updated automatically.

If you are upgrading Data ONTAP nondisruptively (NDU), PAM II firmware is updated nondisruptively. This is because the reboot required for PAM firmware upgrades take place before the final reboot of the `cf giveback` process. Consequently, if your system includes PAM devices, you might see multiple reboots during a Data ONTAP NDU; this is expected behavior.

Firmware updates are not available for the original 16-GB PAM devices. This process refers only to PAM II (256-GB and 512-GB) devices.

For information about what PAM is and how it works, see the *Data ONTAP 7-Mode System Administration Guide*.

Reversion to a previous release

When you create a back-out plan for your Data ONTAP upgrade, you should review reversion guidelines and notices to familiarize yourself with issues you might need to resolve if a reversion becomes necessary. You should contact technical support if you need to revert to a previous release of Data ONTAP.

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

Note: The Upgrade Advisor will identify potential reversion issues when you use it to help you create a back-out plan.

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 Data ONTAP 7.2 family from a release in the 8.0 family, you must review and resolve reversion issues associated with the 7.2 and 7.3 release families before reverting.

In some cases, you cannot revert to an earlier version of Data ONTAP.

Next topics

[General guidelines for reverting from the Data ONTAP 8.0 release family](#) on page 118

[Guidelines for reverting systems with SnapMirror enabled](#) on page 119

[Issues when reverting from Data ONTAP 8.0](#) on page 120

[Issues when reverting to Data ONTAP 7.2](#) on page 126

Related tasks

[Using the Upgrade Advisor to plan your upgrade](#) on page 19

General guidelines for reverting from the Data ONTAP 8.0 release family

You must follow some guidelines before you revert to a previous Data ONTAP version.

Note: Before starting the revert process, ensure that you have a current Snapshot copy of the root volume of any system being reverted. For more information about creating Snapshot copies, see the *Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide*.

The following guidelines apply when you plan to revert from the 8.0 release family to an earlier version:

- You must disable any 8.0 release family features before reverting.

Note: You cannot revert a Data ONTAP 8.0 system if it contains 64-bit aggregates. If your system contains 64-bit aggregates, you must migrate their data to 32-bit aggregates and destroy the 64-bit aggregates before reverting.

- You cannot revert to a release earlier than Data ONTAP 7.2.
- In some cases, you cannot revert to a Data ONTAP release earlier than the one that initially shipped with your system.

If you need functionality from an earlier Data ONTAP release, contact your NetApp representative.

- If you added hardware components after upgrading from an earlier Data ONTAP release, you must verify that the components will continue to work when you revert to the earlier release.

Note: If you upgraded Data ONTAP for new hardware support, you must disconnect the new hardware and reconfigure your system before reverting.

- You cannot revert if an upgrade is in progress. You must complete the upgrade before reverting.
- Before reverting to an earlier release family, you must delete any Snapshot copies made on Data ONTAP release families later than the target release.
- In some cases, the file system identifiers (FSIDs) of volumes on your storage system are rewritten during a revert to be compatible with the version to which you are reverting. Volumes with FSIDs that were rewritten need to be remounted.
- FlexVol volumes must be online before reverting.

If you are reverting to an earlier Data ONTAP release that supports FlexVol volumes, you cannot complete the reversion if there are FlexVol volumes in an offline or restricted state. You must bring these volumes online or destroy them before continuing with the reversion process.

Note: Space guarantees are honored only for online volumes. If you take a volume offline, any committed but unused space for that volume becomes available for other volumes in that aggregate. When you bring that volume back online, there might not be sufficient available space in the aggregate to fulfill its space guarantees.

For more information about space guarantees, see the *Data ONTAP 7-Mode Storage Management Guide*.

- Space guarantees do not persist through reversions to earlier Data ONTAP software versions that support FlexVol volumes.

When you revert to an earlier release, writes to a specified FlexVol volume or writes to files with space reservations enabled could fail if there is not sufficient space in the aggregate.

For more information about space guarantees, see the *Data ONTAP 7-Mode Storage Management Guide*.

- You cannot revert if the background quota upgrade is still in process from a previous Data ONTAP upgrade.

Before reverting to an earlier release, you must either disable quotas or allow the quota upgrade to complete. You can use the `quota status` command to determine if the quota upgrade is still in process.

Guidelines for reverting systems with SnapMirror enabled

If you have enabled SnapMirror data protection on your systems, there are issues to resolve before reverting.

Next topics

[Order for SnapMirror system reversions](#) on page 119

[Preservation of SnapMirror relationships after reversion](#) on page 119

Order for SnapMirror system reversions

If you are reverting on storage systems that are running SnapMirror software, you must revert the systems that have SnapMirror source volumes before you revert the 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 a storage system with SnapMirror source volumes, 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 system, even if the destination system supports that feature.

Preservation of SnapMirror relationships after reversion

During a revert operation, all the Snapshot copies created by the newer version of Data ONTAP are deleted. By performing certain tasks for the Snapshot copy on the source before you upgrade to the newer version of Data ONTAP, you can preserve SnapMirror relationships if you need to revert.

After upgrading Data ONTAP, the older SnapMirror Snapshot copies are gradually replaced with the newer Snapshot copies. If you revert to an older version of Data ONTAP after this replacement, there are no Snapshot copies available for the SnapMirror relationship, and the SnapMirror relationship

would need to be initialized again. This means that the initial SnapMirror baseline transfer required for setting up the replication relationship would need to be performed.

To avoid the need to initialize the SnapMirror relationship again after a revert operation, use one of the following options based on whether you use volume or qtree SnapMirror.

Volume SnapMirror	Creating a manual Snapshot copy on the SnapMirror source before upgrading to the newer version of Data ONTAP, and updating the SnapMirror destination with the changes before upgrading to the newer version of Data ONTAP, enable the SnapMirror relationship to continue with incremental updates, after a revert operation. The manually created Snapshot copy enables you to restore the SnapMirror relationship.
Qtree SnapMirror	Renaming the common Snapshot copy for the qtree SnapMirror relationship on the SnapMirror source before upgrading to the newer version of Data ONTAP, and updating the SnapMirror destination with the changes before upgrading to the newer version of Data ONTAP, enable the SnapMirror relationship to continue with the incremental updates, after a revert operation. The renamed Snapshot copy enables you to restore the SnapMirror relationship.

Attention: After the upgrade, use discretion when deleting any of the older Snapshot copies. After you are sure that a revert operation is not required, you can delete the Snapshot copies from the older version of Data ONTAP.

Issues when reverting from Data ONTAP 8.0

You must understand and resolve issues before you revert from a Data ONTAP 8.0 release.

Next topics

[Disabling compression for SnapMirror transfers after downgrading to Data ONTAP 8.0](#) on page 121

[Reinstatement of in-order frame delivery after reversion](#) on page 121

[Requirements for reverting a system with SSDs attached](#) on page 121

[Retention of modified security settings](#) on page 121

[Reversion issues for Brocade switches in fabric-attached MetroCluster](#) on page 122

[Changes to the interface group configuration in the /etc/rc file](#) on page 122

[Reverting with VLANs and an IP address configured on the base interface](#) on page 122

[Enabling TOE after reverting from Data ONTAP 8.0](#) on page 123

[Downgrade of deduplicated volumes with increased maximum size to Data ONTAP 8.0](#) on page 123

[Reversion of deduplicated volumes with increased maximum size](#) on page 124

[Reverting a SnapMirror destination system with volumes that use deduplication or clone operations](#) on page 124

[Reverting systems when a FlexClone file or FlexClone LUN operation is in progress](#) on page 124

Reverting when Kerberos Multi Realm support is enabled on page 125

Disabling compression for SnapMirror transfers after downgrading to Data ONTAP 8.0

If you enabled compression for SnapMirror transfers, you must disable the feature after you downgrade to Data ONTAP 8.0. Otherwise, you might experience unexpected behavior.

Steps

1. From the SnapMirror destination storage system, open the `/etc/snapmirror.conf` file.
2. Remove the `compression=enable` option from the following entry to disable compression for SnapMirror transfers:

```
connection_name:src_vol dst_system:dst_vol compression=enable * * * *
```

After removing the `compression=enable` option, the entry looks like the following:

```
connection_name:src_vol dst_system:dst_vol - * * * *
```

Reinstatement of in-order frame delivery after reversion

Because out-of-order frame delivery is introduced in Data ONTAP 8.0.1, if you enabled it and then revert to any earlier Data ONTAP release, this functionality is disabled. You must manually enable the in-order delivery options and port-based policy on FC switches.

For more information about enabling in-order frame delivery, see the *Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide* and your FC switch documentation.

Requirements for reverting a system with SSDs attached

SSDs are not supported for any Data ONTAP release earlier than 8.0.1. If you have added SSDs to your system, you must destroy any aggregate made up of SSDs and remove the SSDs from your system before reverting to any earlier version of Data ONTAP.

If you want to preserve the data in the SSDs aggregate, you can use a replication technology such as SnapMirror to copy the data to another aggregate, or you can physically move the SSD aggregate to another system running Data ONTAP 8.0.1.

For more information about SSDs and working with aggregates, see the *Data ONTAP 7-Mode Storage Management Guide*. For more information about SnapMirror and other data replication technologies, see the *Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide*.

Retention of modified security settings

If you upgrade to Data ONTAP 8.0 and subsequently modify your security settings, the modified security settings remain intact even if you later revert to an earlier release of Data ONTAP.

On storage systems shipped with Data ONTAP 8.0 or later, secure protocols are enabled and nonsecure protocols are disabled by default. If you upgrade from an earlier release, existing security settings are not changed to conform to the new defaults. However, if you modify the security settings

after the upgrade and later revert to an earlier release of Data ONTAP, the modified security settings remain intact and the original pre-upgrade security settings are not restored.

Reversion issues for Brocade switches in fabric-attached MetroCluster

Before reverting a fabric-attached MetroCluster to a release prior to Data ONTAP 8.0.1, you must ensure that the Brocade switches 300 or 5100 running on Brocade Fabric OS 6.3.1c or later are downgraded to Brocade Fabric OS supported by that Data ONTAP release.

For more information about the version compatibility of the Brocade Fabric OS with Data ONTAP, see the MetroCluster Compatibility Matrix available on the NOW site.

If the switches are running Brocade Fabric OS 6.3.1c or later, you must first downgrade the switches to Fabric OS 6.2.x before downgrading to Fabric OS 6.1.1a.

Optionally, you can remove the single loop zones from both primary and secondary switches when you are no longer running Fabric OS 6.3.1c or later on them.

Changes to the interface group configuration in the `/etc/rc` file

If you revert from Data ONTAP 8.0 to the Data ONTAP 7.3 or 7.2 release family, the `ifgrp` command entries in the `/etc/rc` file are automatically replaced with `vif` command entries.

Reverting with VLANs and an IP address configured on the base interface

If you configured an IP address for the base interface and configured VLANs on that interface, you must remove the base interface configuration from the `/etc/rc` file before reverting. IP address configuration on the base interface is not supported in the Data ONTAP 7.2 and 7.3 release families.

Steps

1. Open the `/etc/rc` file in the root volume by using a text editor.
2. Delete the command for configuring the IP address on the base interface from the `/etc/rc` file.

Note: If you need the IP address of the base interface after reverting, you can configure it on a different interface.

Example

```
ifconfig e0a 192.0.2.21 netmask 255.255.255.0
vlan create e0a 10 20 30
ifconfig e0a-10 192.0.2.18
ifconfig e0a-20 192.0.2.19
ifconfig e0a-30 192.0.2.20
```

3. To ensure that the VLANs are created successfully after reverting, verify that the command to create the VLANs is listed first, followed by the commands to configure the VLANs.

Example

```
vlan create e0a 10 20 30
ifconfig e0a-10 192.0.2.18
```

```
ifconfig e0a-20 192.0.2.19
ifconfig e0a-30 192.0.2.20
```

Enabling TOE after reverting from Data ONTAP 8.0

Because TOE is automatically disabled in Data ONTAP 8.0, you must enable it either by contacting technical support (for reversion to Data ONTAP 7.3.2 and later) or by modifying your protocol-based options (for reversion to releases earlier than Data ONTAP 7.3.2).

Before you begin

Your system must be running Data ONTAP 7.3.1 or earlier (after you revert from Data ONTAP 8.0) to enable TOE manually. If your system is running Data ONTAP 7.3.2 or later, you must contact technical support to reenable TOE.

Steps

1. To enable TOE in releases earlier than Data ONTAP 7.3.2, enter the following command:

```
options ip.tcp.offload.protocol.enable on
```

2. To enable TOE over protocols, enter the following command:

```
options ip.tcp.offload.protocol.protocol_type on
```

protocol_type can be *iscsi*, *cifs*, or *nfs*.

For more information about the protocol-based options for TOE in a release, see the `na_options (1)` man page for that release.

Downgrade of deduplicated volumes with increased maximum size to Data ONTAP 8.0

Data ONTAP 8.0.1 and later releases supports larger maximum size values for deduplicated volumes. However, if you have increased the size of any deduplicated volume beyond the volume size that is supported in the Data ONTAP 8.0 release, then deduplication will be disabled on that volume when the system boots with ONTAP 8.0.

Checkpoints are deleted during the next deduplication run.

To prevent this, you should either run the `sis undo` command before downgrading or bring the deduplicated volumes to Data ONTAP 8.0 limits.

Note: If you try to enable deduplication on the volume without bringing the deduplicated volume to Data ONTAP 8.0 limits, the deduplication metadata is lost.

Reversion of deduplicated volumes with increased maximum size

Various problems might result if you have increased the size of any deduplicated volume beyond the size supported in an earlier Data ONTAP release and you want to revert to that earlier release.

If the size of any deduplicated volume is beyond the volume size that is supported in Data ONTAP 7.3.1 or later releases, then deduplication is disabled on that volume when the system boots with Data ONTAP 7.3.1 or later.

To prevent this, you should either run the `sis undo` command before downgrading or bring the deduplicated volumes to Data ONTAP 7.3.1 or later limits.

However, if you have increased the size of any deduplicated volume beyond the volume size that is supported in a release earlier than Data ONTAP 7.3.1, then that volume goes offline when the system boots with that earlier release.

Note: If you enable deduplication on the volume without decreasing the volume limits, the deduplication metadata is lost.

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

Reverting a SnapMirror destination system with volumes that use deduplication or clone operations

For a volume SnapMirror relationship, the destination storage system should use an identical or later release of Data ONTAP than the source system.

In releases prior to Data ONTAP 7.3.1, when replicating volumes with deduplication, the NearStore personality license was required on the destination system. However, for Data ONTAP 7.3.1 and later releases, it is not essential to enable the NearStore personality license on the destination system for replicating such volumes. Therefore, if you revert from Data ONTAP 7.3.1 or later to a prior release, you should ensure that the NearStore personality license is enabled on the destination system. Otherwise, after the revert operation, volume SnapMirror updates fail for any volumes on the source that use deduplication.

Note: When using SnapMirror to replicate volumes that use deduplication or clone operations, the destination system should support deduplication.

For more information about the NearStore personality license and the storage systems that support deduplication, see the *Data ONTAP 7-Mode Storage Management Guide*.

Reverting systems when a FlexClone file or FlexClone LUN operation is in progress

Starting with Data ONTAP 7.3.1, you can clone files and LUNs in a FlexVol volume using the FlexClone technology. If you are using FlexClone technology and want to revert to a release earlier

than Data ONTAP 7.3, you should ensure that no FlexClone file or FlexClone LUN operations are in progress.

If any cloning operation is in progress, the presence of temporary Snapshot copies which are used by FlexClone file and LUN operation causes the revert process to fail.

Note: In Data ONTAP 7.3.1 the commands related to FlexClone files and LUNs are available in the `priv set` advanced mode.

For more information about FlexClone volumes, FlexClone files and LUNs, see the *Data ONTAP 7-Mode Storage Management Guide*.

Reverting when Kerberos Multi Realm support is enabled

In Data ONTAP 7.3.1 and later releases, you can configure Data ONTAP to use both Active Directory and UNIX-based KDC types simultaneously. This configuration is sometimes referred to as a Kerberos Multi Realm configuration. However, if you have enabled Multi Realm support on your system and want to revert to an earlier Data ONTAP release, you must take steps *before* reverting.

For more information, see the section on Kerberos security services in the *Data ONTAP 7-Mode File Access and Protocols Management Guide*.

Next topics

[Downgrading to Data ONTAP 7.3 when Kerberos Multi Realm support is enabled](#) on page 125

[Reverting to an earlier release family when Kerberos Multi Realm support is enabled](#) on page 126

Downgrading to Data ONTAP 7.3 when Kerberos Multi Realm support is enabled

If you enable Kerberos Multi Realm support in Data ONTAP 7.3.1 or later and then downgrade to Data ONTAP 7.3, you must first disable Kerberos authentication for NFS. If you reenable Kerberos authentication for NFS after the downgrade and you want to reuse your UNIX keytab file, you must rename the keytab file from `/etc/UNIX_krb5.keytab` to `/etc/krb5.keytab`.

Steps

1. Disable Kerberos for NFS by entering `nfs setup` and answering `y` at the prompt.

```
tpubs-ex1> nfs setup
Kerberos is presently enabled for NFS.
Disable Kerberos for NFS? y
Kerberos now disabled for NFS.
NFS setup complete.
```

2. Downgrade from Data ONTAP 7.3.1 to Data ONTAP 7.3.
3. If you reenable Kerberos authentication for NFS after the downgrade and you want to reuse your UNIX keytab file, you must rename the keytab file from `/etc/UNIX_krb5.keytab` to `/etc/krb5.keytab`.

Note: If you reenables Kerberos authentication for NFS for Data ONTAP 7.3 and later decide to upgrade again to 7.3.1, you must rename the keytab file from `/etc/krb5.keytab` to `/etc/UNIX_krb5.keytab` after upgrading, even if you do not run the `nfs setup` command.

Reverting to an earlier release family when Kerberos Multi Realm support is enabled

If you enable Kerberos Multi Realm support in Data ONTAP 7.3.1 or later and then revert to a Data ONTAP release earlier than 7.3, Data ONTAP automatically disables Kerberos for NFS. You can reenables Kerberos for NFS after such a reversion by running the `nfs setup` command.

Steps

1. Revert Data ONTAP 7.3.1 or later to a Data ONTAP release earlier than 7.3.

If Kerberos Multi Realm support was enabled in Data ONTAP 7.3.1 or later, Data ONTAP displays the following message:

```
*****
Kerberos for NFS will be disabled. If you wish to run
Kerberos for NFS on the reverted release, you need to run
"nfs setup" after revert. If the configuration being used
for NFS after revert will be the same as at present, the NFS
keytab file /etc/UNIX_krb5.keytab can be reused
by renaming it to /etc/krb5.keytab.
*****
```

2. To reenables Kerberos for NFS (and disable Kerberos for CIFS) after the reversion, enter the following command:

```
nfs setup
```

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

3. To reuse your UNIX keytab file, rename it from `/etc/UNIX_krb5.keytab` to `/etc/krb5.keytab`.

Issues when reverting to Data ONTAP 7.2

You must understand and resolve issues before you revert to the Data ONTAP 7.2 release family.

Next topics

[FlexCache reversion limitations](#) on page 127

[Deduplication reversion limitations](#) on page 127

[SnapMirror and SnapVault restart checkpoints deleted during reversion](#) on page 128

[SnapVault licenses might need to be removed before reverting](#) on page 128

[SnapVault restore processes must be complete before reverting](#) on page 128

[Large NFSv4 ACLs removed when reverting from Data ONTAP 7.3](#) on page 128

FPolicy reversion issue with file names having long extensions on page 129

FlexCache reversion limitations

If your storage system has FlexCache volumes, you must ensure that the appropriate license is installed after reverting to a supported release.

FlexCache configurations in Data ONTAP 7.3 and later releases require the new flexcache_nfs license. If you revert to Data ONTAP 7.2.4 or later, the new flexcache_nfs license remains valid. However, if you revert to a Data ONTAP release earlier than 7.2.4, you will not be able to access data in FlexCache volumes unless the old flex_cache license is enabled.

The following Data ONTAP releases support FlexCache volumes:

- 7.3 release family; all releases
- 7.2 release family; Data ONTAP 7.2.1 and later (7.2.4 and later is recommended)

Attention: If you need FlexCache, do not revert to any release earlier than 7.2.1 in the 7.2 release family. You will not be able to access data on FlexCache volumes after reverting to any of these releases.

Deduplication reversion limitations

If you use deduplication, you must ensure that there is adequate free space in the deduplicated volumes and reenable deduplication after reverting to the Data ONTAP 7.2 release family. You also need to run the `sis start -s` command on each deduplicated volume. If you use deduplication, you must ensure that there is adequate free space in the deduplicated volumes.

In Data ONTAP 7.3 and later releases, the deduplication fingerprint database for a volume is stored in the containing aggregate. When you revert to an earlier release, you must ensure that the volume has free space equal to at least 6 percent of its data usage. This space allows the fingerprint database to be recreated in the volume. If sufficient space is not available in the volume, you cannot deduplicate new blocks with ones that existed before the reversion.

For example, if a FlexVol volume has 5 TB of total data (1 TB used and 4 TB saved as reported by the `df -s` command), 300 GB (6 percent of 5 TB) must be available after the revert.

For all Data ONTAP releases earlier than 7.3, after ensuring that this space is available, you must run the `sis start -s` command on every deduplicated volume to rebuild the fingerprint database and reenable deduplication. System resource availability should be considered when determining how many simultaneous deduplication processes to run.

In Data ONTAP 7.3 and later releases, the deduplication fingerprint database for a volume is stored in the containing aggregate. When you revert to an earlier release, you must ensure that the volume has free space equal to at least 6 percent of its data usage. This space allows the fingerprint database to be recreated in the volume.

When you revert from Data ONTAP 8.0 to an earlier release, the system prompts you to run the `sis revert_to` command. The `sis revert_to` command reverts the metadata to the specified earlier release.

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

SnapMirror and SnapVault restart checkpoints deleted during reversion

Starting with Data ONTAP 7.3, when you revert to an earlier version, all aborted qtree SnapMirror and SnapVault transfers with restart checkpoints restart from the beginning because all restart checkpoints are deleted during the reversion process.

SnapVault licenses might need to be removed before reverting

Beginning with Data ONTAP 7.3, you can enable both primary and secondary licenses for SnapVault on the same high availability system node or on a single-node system. However, if both primary and secondary licenses are installed on the same storage system, SnapVault stops functioning after reverting to a Data ONTAP release earlier than 7.3. To continue using SnapVault, you must remove one of the two licenses before reverting.

SnapVault restore processes must be complete before reverting

If you are running SnapVault on Data ONTAP 7.3 or later, you must ensure that any SnapVault restore process has completed before reverting to a Data ONTAP release earlier than 7.3. If a SnapVault restore process is detected, the revert operation will not be allowed to proceed.

You can ensure that no restore process is running by using the following command on the SnapVault secondary storage system:

```
snapvault abort -h [dst_system:]dst_path
```

You must execute this command with the `dst_path` argument for each SnapVault process that is in progress (pending).

The `snapvault abort` process should be allowed to complete before initiating the revert procedure.

Note: The SnapVault restore process cannot be restarted after reverting. If an ongoing SnapVault restore process is critical, allow it to complete before initiating the revert process.

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

Large NFSv4 ACLs removed when reverting from Data ONTAP 7.3

Beginning with Data ONTAP 7.3, the maximum number of Access Control Entries (ACEs) in an Access Control List (ACL) is increased from 192 to 400. If you revert to a release earlier than 7.3, any NFSv4 ACLs with more than 192 ACEs are removed. Files and directories that were created with any of the large ACLs do not have their permissions changed (mode bits are preserved).

Note: If you revert a SnapMirror source system where large NFSv4 ACLs were set on mirrored files or directories, the corresponding files and directories on the destination system will have restrictive ACLs set, which allow only the owner to access them. For more information about reverting SnapMirror systems, see the general guidelines for reverting.

FPolicy reversion issue with file names having long extensions

Starting with Data ONTAP 7.3, the file name extension length supported by FPolicy for file screening is increased to 260 characters. However, if you added longer extensions to the list of extensions to be screened in Data ONTAP 7.3 and you then revert to an earlier version, the file names with the long extensions are not screened by FPolicy after reverting. You should check your FPolicy extension list before reverting.

Optimal service availability during upgrades

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

Next topics

[How upgrades impact service availability](#) on page 131

[Service and protocol considerations](#) on page 132

How upgrades impact service availability

You can review the factors that can affect the availability of storage system services before you begin the upgrade.

The following factors impact service availability:

- Whether the systems being upgraded (upgrade host) are single nodes or HA configuration partners
Systems in high-availability configurations are designed to provide optimal service availability.
- The types of protocols used and services licensed, and their susceptibility to timeout errors
- Whether you need to make decisions about Data ONTAP issues and new features between or within release families
Upgrading between Data ONTAP release families involves more steps and is potentially more disruptive than upgrades within a release family.
- Whether a system firmware update is required
Some system firmware updates require a system halt and reboot. This can disrupt services in single systems and standard HA configuration upgrades but does not affect services in nondisruptive HA configuration upgrades.
- Whether a disk shelf firmware update is required
Nondisruptive firmware upgrades are available for many disk shelf and module configurations.
- Whether disk firmware updates are required and what type of RAID protection applies to those disks
- 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.

Related concepts

[Upgrade host requirements](#) on page 23

[Service and protocol considerations](#) on page 132

[Evaluating upgrade issues](#) on page 32

[Updating firmware](#) on page 87

[Disk shelf firmware updates](#) on page 97

[Disk firmware updates](#) on page 92

Service and protocol considerations

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.

During an upgrade, the storage system must be rebooted (by issuing the `reboot` command or by initiating an HA configuration takeover and giveback) to load the new software. Services based on stateless protocols usually remain available during nondisruptive upgrades of systems in an HA configuration.

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 HA configuration 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.

Next topics

[Considerations for stateless protocols](#) on page 132

[Considerations for session-oriented protocols](#) on page 133

Considerations for stateless protocols

Configurations that include client connections using stateless protocols generally do not experience adverse effects during upgrade if the clients are configured according to recommended guidelines.

- NFS hard mounts
No adverse behavior on the clients. Clients might receive some messages similar to the following until the storage system reboots:
`NFS server not responding, retrying`
In general, read/write directories should be hard mounted. Hard mounts are the default type of mount.
- NFS soft mounts
You should not use soft mounts when there is a possibility of frequent NFS timeouts. Race conditions can occur as a result of these timeouts, which can lead to data corruption. Furthermore,

some applications cannot properly handle errors that occur when a NFS operation reaches a timeout using soft mounts.

Some of the situations that can cause frequent timeouts are nondisruptive upgrades or any takeover/giveback event in an HA configuration.

In general, soft mounts should be used only when solely reading from a disk. Even then, understand that the mount is unreliable.

- SAN protocols

No adverse behavior on FC or iSCSI clients provided they are configured according to recommended guidelines.

For more information, see the *NetApp Interoperability Matrix* on the NOW site.

Related information

[NetApp Interoperability Matrix -- now.netapp.com/NOW/products/interoperability/](http://now.netapp.com/NOW/products/interoperability/)

Considerations for session-oriented protocols

Storage systems and session-oriented protocols might cause adverse effects on clients and applications in the following areas during upgrades.

- CIFS

Client sessions are terminated. You should inform users to end their sessions before you upgrade. To do so, issue the following command before the HA configuration takeover:

```
cifs terminate -t
```

Alternatively, issue the following command before the reboot:

```
reboot -t
```

- 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 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

- A**
- Active Directory-based KDC 41
- B**
- BMC firmware 111
- C**
- CIFS 29
 - requires standard upgrade 29
 - CPU utilization, nondisruptive upgrade requirements 30
- D**
- Data ONTAP 43, 45, 73, 78, 81, 84, 117, 118, 120, 123, 124, 126
 - downgrading deduplicated volumes 123
 - guidelines for reverting from the 8.0 release family 118
 - preparing for the upgrade 43
 - reverting deduplicated volumes 124
 - reverting from Data ONTAP 8.0 120
 - reverting to a previous release 117
 - reverting to Data ONTAP 7.2 126
 - upgrading a high-availability configuration (standard) 81
 - upgrading a high-availability configuration from an earlier release family nondisruptively 73
 - upgrading a high-availability configuration within a release family (nondisruptive) 78
 - upgrading a single system 84
 - version supported 45
 - Data ONTAP 7.3.1 and later 41
 - Kerberos Multi Realm support 41
 - Data ONTAP 8.0 122
 - reversion issues with VLANs 122
 - Data ONTAP software images 56, 57, 59, 60
 - copy software images without installing 56
 - copying from a UNIX client 57
 - copying from a Windows client 59
 - getting from NOW 60
 - Data ONTAP system files 56, 58, 61, 63, 64, 67
 - copying the software image to the HTTP server 56
 - downloading from NOW 58
 - installation overview 63
 - installation procedure for HTTP 64
 - installation procedure from /etc/software 67
 - managing from an HTTP server 56
 - managing with the software command 61
 - deduplication 46
 - upgrade requirements 46
 - disk firmware upgrades 92, 94, 95
 - about 92
 - background 94
 - standard 95
 - disk shelf firmware upgrades 97–100
 - about 97
 - determining firmware versions 99
 - manual update procedure 100
 - service availability during 98
 - disk utilization, nondisruptive upgrade requirements 30
 - disk_fw_update command 96
 - DNS, enable 47
 - domain account, verifying 47
 - downgrade issues 121
 - compression for SnapMirror transfers 121
 - downgrading from Data ONTAP 8.0 122
 - VLAN configuration issues 122
- F**
- firmware upgrades 87, 92, 97, 103, 105, 111, 116
 - BMC 111
 - disk 92
 - disk shelf 97
 - PAM 116
 - RLM 105
 - SP (Service Processor) 103
 - system 87
 - FlexCache 35
 - origin system version requirement 35
 - FlexVol volumes 30
 - nondisruptive upgrade requirements 30
- H**
- HA Configuration Checker 47

I

in-order frame delivery, reverting with 121

K

Kerberos, Multi Realm support 41, 125

L

LUN restore 37

M

major 30
 nondisruptive upgrades 30
 minor 30
 nondisruptive upgrades 30
 module firmware, disk shelf 97
 Multi Realm support, Kerberos 41, 125

N

nondisruptive upgrades 28–30, 48, 88, 100
 about 28
 Data ONTAP software 48
 disk shelf firmware, not supported 100
 preparing 48
 requirements 30
 system firmware 88
 using Upgrade Advisor tool 30
 when not to use 29

P

PAM firmware 116

R

release families 26, 27
 differentiating among 26
 overview 26
 upgrading between 27
 upgrading within 27
 reversion issues 122, 125
 Compatibility issues attached with Brocade Fabric
 Operating System and Data ONTAP 122
 Kerberos Multi Realm support 125
 reversion issues FlexClone files and LUNs, reverting
 124
 FlexClone files and LUNs 124

revert 119
 SnapMirror, preserve relationship 119
 reverting from Data ONTAP 8.0 122, 123
 enabling TOE 123
 VLAN and base interface configuration issues 122
 reverting to a previous release 117, 118, 120, 126
 reverting from Data ONTAP 8.0 120
 reverting from the 8.0 release family 118
 reverting to Data ONTAP 7.2 126
 technical support 117
 revision issues 121
 in-order frame delivery 121
 RLM 110
 firmware update problems, troubleshooting 110
 troubleshooting firmware update problems 110
 RLM firmware 105
 rolling upgrade 28

S

shelf, disk 97
 SnapLock 33, 34
 considerations before upgrading 33, 34
 SnapMirror 24–26, 71, 119
 identifying destination volumes 71
 issues for systems with synchronous SnapMirror 25
 planning upgrades 25
 revert, initialize 119
 revert, preserve relationship 119
 upgrade requirements 24
 upgrade, preserve relationship 119
 upgrading for volume replication 71
 upgrading systems that are mirroring volumes to
 each other 26
 Snapshot copies 30
 nondisruptive upgrade requirements 30
 software update command 63
 solid-state disks (SSDs) 121
 reverting and 121
 SP (Service Processor) firmware 103
 special system files 37
 .bplustvoc_internal 37
 .vtoc_internal 37
 SSDs 121
 reverting and 121
 standard system firmware update 91
 storage download shelf command 100
 system firmware 87, 88, 91
 about 87
 nondisruptive upgrade 88

- obtaining 88
- standard firmware update procedure 91

U

- UNIX host 58
 - mounting the system 58
- UNIX-based KDC 41
- upgrade 19, 21, 28, 32, 43, 44, 47, 119, 131–133
 - enabling DNS with Windows 2000 name addresses 47
 - maintaining service 131
 - overview 21
 - overview of requirements 19
 - planning 19
 - preparing for 43
 - required intermediate upgrades 28
 - resolving issues 32
 - SnapMirror, preserve relationship 119

- system requirements 44
- verifying system domain account 47
 - with session oriented protocols 133
 - with stateless protocols 132
- Upgrade Advisor tool 19, 30
 - about 19
- upgrade host 23
 - requirements 23

V

- VLAN configuration 122
 - reversion issues with Data ONTAP 8.0 122

W

- Windows host 59
 - mapping the root directory to a client share 59

