

IBM System Storage N series



Data ONTAP 8.0 7-Mode Upgrade and Revert/Downgrade Guide

Contents

Preface	9
About this guide	9
Supported features	9
Websites	9
Getting information, help, and service	10
Before you call	10
Using the documentation	10
Hardware service and support	11
Firmware updates	11
How to send your comments	11
Planning your upgrade	13
Upgrade process overview	13
Recommendations for all systems upgrading to this release	15
Upgrade host requirements	15
Requirements when upgrading from a Windows or UNIX client using the CIFS or NFS protocols	16
Requirements when upgrading from an HTTP server	16
Upgrade requirements for SnapMirror	16
Why you must plan for SnapMirror upgrades	17
SnapMirror synchronous and asynchronous mode during upgrade	17
Upgrade requirements for systems mirroring each other	18
Release family upgrade requirements	18
Different types of upgrades	18
Upgrades between release families	18
Upgrades within a release family	19
Required intermediate upgrades	19
Nondisruptive upgrade requirements	20
When to use nondisruptive high-availability upgrades	20
When not to use nondisruptive upgrades	20
Requirements for nondisruptive upgrades on all systems	21
Requirements for nondisruptive upgrades on systems with deduplicated volumes	22

Disruptive upgrade requirements	23
Evaluating upgrade issues	23
Issues to resolve before upgrading to the Data ONTAP 8.0 release family	24
Changes to behavior in the Data ONTAP 8.0 release family	25
Issues to resolve before upgrading from releases earlier than Data ONTAP 7.3	25
Behavior changes when upgrading from releases earlier than Data ONTAP 7.3	26
Preparing for the upgrade	27
Verifying system requirements	28
Ensuring that there is adequate free space in every volume containing LUNs	28
Deduplication upgrade requirements	28
Determining the required firmware for your disks	28
Determining the required firmware for your disk shelves	28
Enabling DNS with Windows 2000 name server addresses	29
Verifying that you have a domain account	29
Preparing for nondisruptive upgrades	29
Preparing for nondisruptive upgrades on systems with VMware ESX server hosts	32
Determining system capacity and space guarantees before upgrading to Data ONTAP 7.3 or later	33
Using the aggrSpaceCheck tool to prepare your upgrade to Data ONTAP 7.3 or later	34
Reconfiguring IPv4 before upgrading	35
Obtaining Data ONTAP software images	37
Obtaining images for HTTP servers	37
Copying the software image to the HTTP server	38
Copying software images from the HTTP server without installing the images	38
Obtaining images for UNIX clients	39
Mounting the storage system on your client	39
Obtaining software images	39
Obtaining images for Windows clients	40
Mapping the storage system to a drive	41

Obtaining software images	41
Managing files in the /etc/software directory	42
Installing Data ONTAP software images	43
Installing software images from an HTTP server	43
Installing software images from the /etc/software directory	47
Downloading and rebooting new Data ONTAP software	51
Upgrading in a SnapMirror environment	51
Upgrading nondisruptively in a SnapMirror environment	52
Upgrading HA configurations from an earlier release family nondisruptively	53
Upgrading HA configurations within a release family nondisruptively	58
Upgrading HA configurations using the disruptive method	60
Upgrading single systems	63
Verifying system status after an upgrade	65
Updating IBM customer contact information	67
Entering customer contact information with the setup command	67
Updating firmware	69
System firmware updates	69
Automatic BIOS system firmware updates	69
Updating system firmware nondisruptively	70
Updating system firmware using the disruptive method	72
Disk firmware updates	73
How disk firmware is updated	73
Service availability during disk firmware updates	74
Updating disk shelf firmware	75
How disk shelf firmware is updated	75
Disk shelf firmware requirements for Data ONTAP nondisruptive upgrades	76
Detecting outdated disk shelf firmware	78
Updating disk shelf firmware manually	78
Updating ACP firmware	80
Service Processor firmware updates	81
Using the Data ONTAP CLI to update the SP firmware	81
Using the SP CLI to update the SP firmware	82
RLM firmware updates	82
Requirements for RLM firmware version 4.0 and later	83
Using the Data ONTAP CLI to update the RLM firmware	83

Using the RLM CLI to update the RLM firmware	85
RLM firmware update problems	87
BMC firmware updates	88
Detecting outdated BMC firmware	89
Updating BMC firmware nondisruptively	90
Updating BMC firmware using the disruptive method	92
Flash Cache firmware updates	93
Reverting to an earlier 7-Mode release family	95
When to revert and when to call technical support	95
7-Mode reversion checklist	96
General reversion requirements	96
Requirements for reverting configured systems	98
Special system files	99
Preparing to revert Data ONTAP 7-Mode	99
Commands for addressing reversion requirements	100
Preparing to revert configured systems	103
Staging the target Data ONTAP image	109
Performing the 7-Mode reversion process	109
Reverting Data ONTAP	109
Updating SP firmware	112
Completing post-reversion tasks	113
Using deduplication on a reverted system	113
Reenabling NDMP on a reverted system	114
Enabling TOE after reverting from Data ONTAP 8.0	114
Reinstatement of in-order frame delivery after reversion	114
Downgrading to an earlier release in the same 7-Mode release family	117
When to downgrade and when to call technical support	117
7-Mode downgrade checklist	117
General downgrade requirements	118
Requirements when downgrading to Data ONTAP 8.0 7-Mode	119
Disabling compression for SnapMirror transfers after downgrading to Data ONTAP 8.0	119
Downgrade of deduplicated volumes with increased maximum size to Data ONTAP 8.0	120
Preparing to downgrade Data ONTAP	120
Commands for addressing downgrade requirements	121

Staging the target Data ONTAP image	122
Performing the 7-Mode downgrade process	122
Downgrading Data ONTAP using the nondisruptive method	123
Downgrading Data ONTAP using the disruptive method	125
Updating SP firmware	126
Completing post-downgrade tasks	128
Optimal service availability during upgrades	129
How upgrades impact service availability	129
Service and protocol considerations	130
Considerations for stateless protocols	130
Considerations for session-oriented protocols	131
Understanding background disk firmware updates	131
Copyright information	133
Trademark information	135
Index	137

Preface

About this guide

This guide applies to systems, including systems with gateway functionality, running Data ONTAP 8.x 7-Mode. In the Data ONTAP 8.x 7-Mode product name, the term *7-Mode* signifies that the 8.x release has the same features and functionality found in the prior Data ONTAP 7.1, 7.2, and 7.3 release families.

Note: In this document, the term *gateway* describes IBM N series storage systems that have been ordered with gateway functionality. Gateways support various types of storage, and they are used with third-party disk storage systems—for example, disk storage systems from IBM, HP®, Hitachi Data Systems®, and EMC®. In this case, disk storage for customer data and the RAID controller functionality is provided by the back-end disk storage system. A gateway might also be used with disk storage expansion units specifically designed for the IBM N series models.

The term *filer* describes IBM N series storage systems that either contain internal disk storage or attach to disk storage expansion units specifically designed for the IBM N series storage systems. Filer storage systems do not support using third-party disk storage systems.

Supported features

IBM System Storage N series storage systems are driven by NetApp Data ONTAP software. Some features described in the product software documentation are neither offered nor supported by IBM. Please contact your local IBM representative or reseller for further details.

Information about supported features can also be found on the N series support website, which is accessed and navigated as described in [Websites](#) on page 9.

Websites

IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. The following web pages provide N series information:

- A listing of currently available N series products and features can be found at the following web page:
www.ibm.com/storage/nas/
- The IBM System Storage N series support website requires users to register in order to obtain access to N series support content on the web. To understand how the N series support web

content is organized and navigated, and to access the N series support website, refer to the following publicly accessible web page:

www.ibm.com/storage/support/nseries/

This web page also provides links to AutoSupport information as well as other important N series product resources.

- IBM System Storage N series products attach to a variety of servers and operating systems. To determine the latest supported attachments, go to the IBM N series interoperability matrix at the following web page:

www.ibm.com/systems/storage/network/interophome.html

- For the latest N series hardware product documentation, including planning, installation and setup, and hardware monitoring, service and diagnostics, see the IBM N series Information Center at the following web page:

publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Getting information, help, and service

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your IBM N series product, and whom to call for service, if it is necessary.

Before you call

Before you call, make sure you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure they are connected.
- Check the power switches to make sure the system is turned on.
- Use the troubleshooting information in your system documentation and use the diagnostic tools that come with your system.
- Refer to the IBM N series support website for information on known problems and limitations.

Using the documentation

The latest versions of N series software documentation, including Data ONTAP and other software products, are available on the IBM N series support website, which is accessed and navigated as described in *Websites* on page 9.

Current N series hardware product documentation is shipped with your hardware product in printed documents or as PDF files on a documentation CD. For the latest N series hardware product documentation PDFs, go to the IBM N series support website.

Hardware documentation, including planning, installation and setup, and hardware monitoring, service, and diagnostics, is also provided in an IBM N series Information Center at the following web page:

publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Hardware service and support

You can receive hardware service through IBM Integrated Technology Services. Visit the following web page for support telephone numbers:

www.ibm.com/planetwide/

Firmware updates

IBM N series product firmware is embedded in Data ONTAP. As with all devices, it is recommended that you run the latest level of firmware. Any firmware updates are posted to the IBM N series support website, which is accessed and navigated as described in [Websites](#) on page 9.

Note: If you do not see new firmware updates on the IBM N series support website, you are running the latest level of firmware.

Verify that the latest level of firmware is installed on your machine before contacting IBM for technical support.

How to send your comments

Your feedback helps us to provide the most accurate and high-quality information. If you have comments or suggestions for improving this document, please send them by e-mail to starpubs@us.ibm.com.

Be sure to include the following:

- Exact publication title
- Publication form number (for example, GC26-1234-02)
- Page, table, or illustration numbers
- A detailed description of any information that should be changed

Planning your upgrade

Because new features are introduced in each release of Data ONTAP, you must understand these features and upgrade requirements to 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.

Before proceeding with the upgrade, you should 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.
- Note any potential behavior changes to your system after the upgrade.
- Be prepared to address all points in the upgrade checklist.
- Create a back-out plan, in the unlikely event that you need to revert or downgrade to the Data ONTAP release that was running on your system before the upgrade.

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. If you attempt to download and reboot a Data ONTAP 8.0.x release on a system with SnapLock volumes, the system halts.

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.
 - Create a back-out plan, in the unlikely event that you need to revert or downgrade to the Data ONTAP release that was running on your system before the upgrade.
 - 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 the compatibility and configuration information about FCP and iSCSI products.
See the appropriate matrix at the N series Service and Support website at www.ibm.com/storage/support/nseries/.

- 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 N series support 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 finished when your HA pair or single system reboots with the new version of Data ONTAP.
 6. Verify that your systems are operating as expected after the upgrade.
Before returning storage systems to production, you should check the status of configured functionality and reenable any functionality that was suspended before the upgrade
 7. If you are upgrading from a release earlier than Data ONTAP 7.2.5, supply IBM customer support information at the storage system command-line interface after completing the upgrade.

Related concepts

[Planning your upgrade](#) on page 13

[Updating firmware](#) on page 69

[Obtaining Data ONTAP software images](#) on page 37

[Installing Data ONTAP software images](#) on page 43

[Downloading and rebooting new Data ONTAP software](#) on page 51

[Updating IBM customer contact information](#) on page 67

Related tasks

[Preparing for the upgrade](#) on page 27

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. If you attempt to download or boot a Data ONTAP 8.0.x release on a system with SnapLock volumes, the system halts.

Data ONTAP 8.x operating systems are much larger than in previous release families. When you use the `download` or `software` command to activate Data ONTAP 8.x software images on your storage system boot device, be aware that the `download` process takes significantly longer to finish than for earlier releases. For most systems upgrading to Data ONTAP 8.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 N series support 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

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

Related concepts

[Installing Data ONTAP software images](#) on page 43

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*.

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 37

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*.

Related tasks

[Upgrading in a SnapMirror environment](#) on page 51

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 51

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.

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.

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 21

[Required intermediate upgrades](#) on page 19

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.

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.

For these conditions, disruptive upgrades are recommended.

Related concepts

[Updating disk shelf firmware](#) on page 75

Requirements for nondisruptive upgrades on all systems

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

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' Service Processors, if available

- 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 N3400 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 22

[Optimal service availability during upgrades](#) on page 129

[Considerations for stateless protocols](#) on page 130

[Required intermediate upgrades](#) 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.

Disruptive upgrade requirements

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

In a disruptive 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 disruptive upgrade. These can also take place when the HA pair is disabled.

Although nondisruptive upgrade requirements are not mandatory for disruptive upgrades (when downtime is scheduled), it is a best practice to follow the NDU preparatory procedures for all upgrades to ensure system health before and after the upgrade.

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.

Be sure to consult the *Release Notes* for the upgrade target release for a complete listing of upgrade issues. If an intermediate upgrade is required, you should also consult the *Release Notes* for the intermediate release.

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.

This topic summarizes significant issues known at publication time. Be sure to check the *Important Cautions* section in the *Release Notes* for your target Data ONTAP release to see a complete list of issues that could affect the upgrade.

- **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.
- **IPv6 not supported in Data ONTAP 8.0**
If IPv6 is enabled on your storage system, be aware that it is 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.
Also, 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.
- **IPsec is not supported in Data ONTAP 8.0**
If IPsec is enabled on your storage system, be aware that it is 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.
- **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.

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.

This topic summarizes significant changes known at publication time. Be sure to check the *Known Problems and Limitations* section in the Release Notes for your target Data ONTAP release to see a complete list of changes in behavior after upgrade to the target release.

- **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.
- **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.
- **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.
- **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.

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.

This topic summarizes significant issues known at publication time. Be sure to check the *Important Cautions* section in the *Release Notes* for the intermediate Data ONTAP release to see a complete list of issues that could affect the upgrade to any release later than the intermediate release.

- **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.

- 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.
- Disks offline in Windows 2008 after a disruptive upgrade
During a disruptive 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.

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.

This topic summarizes significant changes known at publication time. Be sure to check the *Known Problems and Limitations* section in the *Release Notes* for the intermediate Data ONTAP release to see a complete list of changes in behavior after upgrade to any release later than the intermediate release.

- The NetBackup application can no longer manage SnapVault relationships with storage system data
Beginning with Data ONTAP 7.3, the use of Symantec NetBackup for configuring and managing SnapVault transfers between primary and secondary storage systems is no longer supported.
- 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.
- 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.
- 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.
- 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.

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.

For more information about system requirements, see the Release Notes for your Data ONTAP upgrade target release.

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. If you are running SnapMirror, identify storage systems with destination volumes and upgrade them before upgrading storage systems with source volumes.
5. 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 gateway, see also the *Gateway Interoperability Matrix*.

6. 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

7. If you are using the nondisruptive upgrade method, ensure that your systems meet the requirements.
8. If you are upgrading from a release earlier than Data ONTAP 7.3, ensure that there is adequate free space in your aggregates.
9. 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.

Related concepts

[Why you must plan for SnapMirror upgrades](#) on page 17

Verifying system requirements

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

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 N series support site, you can determine if you need to update the disk firmware for your system.

Determining the required firmware for your disk shelves

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

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.

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.

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.

2. Ensure that network ports are up and functioning correctly by entering the following command:

```
ifconfig -a
```

Example

For each interface, you see a display similar to the following:

```
e0a: flags=0x2f4c867<UP,BROADCAST,RUNNING,MULTICAST,TCPCKSUM,LINK_UP>
mtu 1500
  inet 192.9.200.41 netmask 0xffffffff broadcast 192.9.200.255
  partner e0a 192.9.200.42
  ether 00:0c:29:56:54:7e (auto-1000t-fd-up) flowcontrol full
```

For each interface that serves data traffic, ensure that each of the following is true:

- a. The interface has a partner that also serves data; that is, the partner is not an e0M or e0P interface.
 - b. The link to the partner is up.
 - c. The `mtu`, `mediatype`, and `flowcontrol` parameter settings are the same for both partners.
3. If you have edited the `/etc/rc` file, ensure that entries are listed in the following order:

```
hostname system_name
ifgrp [commands]
vlan [commands]
ifconfig [commands]
vfiler [commands]
route [commands]
[any other commands]
```

4. If your systems include e0M management interfaces, ensure that they are serving only management traffic on a dedicated management LAN or that they are configured down.

If an e0M interface is serving management traffic, it should be partnered with another e0M interface.

For more information about e0M configuration, see the *Data ONTAP 7-Mode System Administration Guide*.

5. If your systems include e0P interfaces for controlling SAS disk shelves, ensure that they are connected only to a private ACP network or that they are configured down.

e0P interfaces should not be partnered.

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

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

Check service protocols and configure client timeout settings to ensure that they meet the availability requirements for a nondisruptive upgrade.

7. Verify that all components of your SAN configuration are compatible with the upgraded Data ONTAP release by consulting the compatibility and configuration information about FCP and iSCSI products.

See the appropriate matrix at the N series Service and Support website at www.ibm.com/storage/support/nseries/.

8. 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).

9. 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.

10. 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.

11. 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.

Data ONTAP 8.0.2 and later releases support background disk firmware updates for disks attached to non-mirrored RAID4 aggregates. It is no longer necessary to perform disruptive disk firmware updates on systems with these aggregates.

12. 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.

13. If you use SnapMirror technology, ensure that SnapMirror is suspended and no SnapMirror operations are in process while upgrading Data ONTAP.

Related concepts

[Optimal service availability during upgrades](#) on page 129

[Updating disk shelf firmware](#) on page 75

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 IBM N Series Storage Systems* on the N series support 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 IBM Redbook *IBM System Storage N series with VMware ESX Server*.

Related information

IBM System Storage N series with VMware ESX Server: www.redbooks.ibm.com/abstracts/sg247636.html?Open

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 `aggrSpaceCheck` tools to determine your system capacity and plan your upgrade.

Step

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	Use the <code>aggrSpaceCheck</code> tool to plan your upgrade.

After you finish

After using the `aggrSpaceCheck` tool and completing the upgrade, make sure that your space guarantees are configured according to your requirements.

Related tasks

Using the `aggrSpaceCheck` tool to prepare your upgrade to Data ONTAP 7.3 or later on page 34

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.

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 N series support 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 N series support 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. 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.

2. 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.

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 N series support 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 N series support 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.

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.

Related concepts

[Installing Data ONTAP software images](#) on page 43

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 N series support 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 N series support site.

Related concepts

[Upgrade host requirements](#) on page 15

[Installing Data ONTAP software images](#) on page 43

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 N series support site.

Obtaining software images

You can use a web browser to copy the software image from the N series support 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 N series support site.
2. Click **Data ONTAP** from the Current N series NAS/iSCSI product list.
3. Click the **Download** tab, and then click Downloadable Files.
4. Choose **Windows** from the Platform/Operating system drop-down list.
5. 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 N series support 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.

6. 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 N series support site.

Related concepts

Upgrade host requirements on page 15

Installing Data ONTAP software images on page 43

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 N series support site.

Obtaining software images

You can use a web browser to copy the software image from the N series support 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 N series support site.
2. Click **Data ONTAP** from the Current N series NAS/iSCSI product list.
3. Click the **Download** tab, and then click Downloadable Files.
4. Choose **Windows** from the Platform/Operating system drop-down list.
5. 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 N series support 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.

6. 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)` command 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.

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	<code>software update http://www.example.com/downloads/x86-64/my_80_setup.i.tgz -d</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 -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</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 -R</code>

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 51

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 /etc/software 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 /etc/software 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 51

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.

Related concepts

[Release family upgrade requirements](#) on page 18

[Disruptive upgrade requirements](#) on page 23

[Nondisruptive upgrade requirements](#) on page 20

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

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 reenble 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.

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 reenble 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 | At the console of each system, enter the following command:<br><br><b>software update <i>file_name</i> -r</b><br><br>Then go to Step 4.<br><br><b>Note:</b> 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 in earlier releases. The `download` process for Data ONTAP 8.0.x usually takes 20 to 60 minutes.

- 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.

- 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     | Enter the following command:<br><br><b>cifs terminate -t nn</b><br><br><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. |

- 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 cannot boot Data ONTAP 8.0 and the upgrade fails.

- 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
```

- 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 high-availability partner.

**Note:** At this point in the upgrade procedure—system A is running Data ONTAP 7.3 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 not harmful.

Nonetheless, you should complete the upgrade procedure as quickly as possible; do not allow the two systems to remain in a state of version mismatch longer than necessary.

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.

**13.** Choose the following option that describes your configuration.

| <b>If CIFS...</b>         | <b>Then...</b>                                                                                                                                                                                                                                                |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Is not in use in system B | Go to the next step.                                                                                                                                                                                                                                          |
| Is in use in system B     | Enter the following command:<br><br><pre><b>cifs terminate -t nn</b></pre> <p><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.</p> |

**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 generates an error and the command terminates.

**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).

| <b>If...</b>                | <b>Then...</b>                                                   |
|-----------------------------|------------------------------------------------------------------|
| 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 cannot boot Data ONTAP 8.0 and the upgrade fails.

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

**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 immediately terminates 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.

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.

| <b>If you...</b>                                                  | <b>Then...</b>                                                                                                                                                                                                                                                                                                                                                                               |
|-------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Have already installed system files                               | Go to the next step.                                                                                                                                                                                                                                                                                                                                                                         |
| Are installing and downloading system files in the same operation | <p>At the console of each system, enter the following command:</p> <pre><b>software update file_name -r</b></pre> <p>Then go to Step 4.</p> <p><b>Note:</b> 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:** Using the `download` procedure to activate Data ONTAP 8.0 software images takes significantly longer to complete than on earlier releases.

- 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.

- 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:<br><br><pre><b>cifs terminate -t nn</b></pre> <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. |

- 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.

- 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 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 will not affect the nondisruptive upgrade; the final reboot returns the system to high-availability status.

7. Repeat Step 4 through Step 6 to update the partner storage system; in other words, bring down and update system B with partner 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:<br><br><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 disruptive method

If you can take HA pairs offline to update software and other components, you can use the disruptive 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 disruptive 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 | <p>At the console of each system, enter the following command:</p> <pre>software update file_name -r</pre> <p>Then go to Step 4.</p> <p><b>Note:</b> 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 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 47*

## 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><b>software update file_name -r</b></pre> <p>Then go to Step 3.</p> <p><b>Note:</b> 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 47*



## Verifying system status after an upgrade

It is a best practice to verify that upgraded systems are functioning as expected before returning them to production. This entails verifying the status of configured functionality and reenabling any functionality that was suspended before the upgrade.

### About this task

These tasks should be performed on each partner of an HA pair and on every single system that was upgraded.

### Steps

1. If your systems are in an HA configuration, enter the following command to verify that the HA relationship is operational:  
`cf status`
2. If you disabled automatic giveback to prepare for a nondisruptive upgrade and you wish to reenable it, enter the following command:  
`options cf.giveback.auto.enable on`
3. If your systems are configured for MultiStore technology, enter the following command to verify that all vFiler units are running:  
`vfiler status -a`
4. Verify that the intended target release is installed and running by entering the following command:  
`version`
5. Confirm that all paths to disk are visible to the system by entering the following command:  
`storage show disk -p`
6. Confirm that there are no failed disks by entering the following command:  
`vol status -f`
7. Verify that all aggregates are online by entering the following command:  
`aggr status`
8. Confirm that all volumes have a volume guarantee of none (except the root volume, vol0) by entering the following command:  
`vol status -v`
9. Confirm that network interfaces are online by entering the following command:  
`ifconfig -a`

10. If you disabled SnapMirror functionality, enter the following command to reenable it:

```
snapmirror on
```

11. If you quiesced SnapMirror transfers, enter the following command for each destination volume to resume them:

```
snapmirror resume destination
```

12. If you have enabled data compression on any volume, you must check the status of data compression scans. If any were halted by the upgrade, you must restart them on each volume.

| <b>If you want to...</b>                   | <b>Enter this command...</b>                |
|--------------------------------------------|---------------------------------------------|
| Check the status of data compression scans | <code>vol compress status [vol_name]</code> |
| Restart the data compression scanner       | <code>vol compress start vol_name</code>    |

## Updating IBM customer contact information

---

Data ONTAP 7.2.5 and later releases include improved AutoSupport reporting features. To take advantage of these features, you must enter IBM customer contact information after completing the upgrade.

There are two ways to enter IBM customer contact information after upgrading.

- Running the `setup` at the storage system command line.
- Entering values in the customer contact options at the storage system command line.  
For more information about customer contact options, see the *Data ONTAP 7-Mode System Administration Guide*.

Either of these procedures can be used to update customer contact information after initial system setup or upgrade.

## Entering customer contact information with the setup command

You can run the `setup` command after upgrading to this release to enter required IBM customer contact information.

### Before you begin

The upgrade to Data ONTAP 7.2.5 or later must be complete before you update IBM customer contact information.

### About this task

You need to gather contact information and machine location information to enter when completing this task.

For more information, about the `setup` command, see the `setup(1)` man page.

### Steps

1. Record the customer contact information for the upgraded system.

Use the detailed descriptions in the "Required IBM customer information" section of the *Data ONTAP 7-Mode Software Setup Guide* to gather this information.

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

```
setup
```

The setup display describes the files that will be rewritten when you run the command. You will be able to preserve values you have already entered.

3. Enter **y** to continue.

Your current system configuration is displayed (the output of the `sysconfig` command, followed by a series of configuration prompts). The values that you already entered for these parameters are given in square brackets.

4. Press **Enter** to accept the each of the existing values.

Continue until you see the prompts for customer contact information.

5. Enter the contact information you gathered for the following values:

```
Name of primary contact (Required)
Phone number of primary contact (Required)
Alternate phone number of primary contact
Primary Contact e-mail address or IBM WebID
Name of secondary contact
Phone number of secondary contact
Alternate phone number of secondary contact
Secondary Contact e-mail address or IBM WebID
```

6. Enter the machine location you gathered for the following values:

```
Business name (Required)
Business address (Required)
City where business resides (Required)
State where business resides
2-character country code (Required)
Postal code where business resides
```

7. When setup is complete, to transfer the information you've entered to the storage system, enter the following command, as directed by the prompt on the screen.

**reboot**

**Note:** If you do not enter `reboot`, the information you entered does not take effect.

8. If you are configuring a pair of storage systems in an active/active configuration and have not configured the other storage system, repeat these instructions to set up the other storage system in the configuration.

## 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.

### 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 N series support website and installing the files.

If you are upgrading system firmware between Data ONTAP upgrades, you can use the nondisruptive or disruptive methods to update system firmware manually. You can obtain system firmware and information about how to install it from the N series support website.

### 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 8.0.x releases.

| Platform     | Minimum version     |
|--------------|---------------------|
| N7000 series | BIOS 1.5X2 or later |
| N6000 series |                     |

| Platform | Minimum version        |
|----------|------------------------|
| N5600    | BIOS 2.2X1 or later    |
| N5300    |                        |
| N3400    | BIOS/NABL 6.0 or later |

N7x50T series and N6200 series 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 N series support 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 N series support 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. |

| If CIFS...             | Then...                                                                                                                                                                                                                                                               |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Is in use in system A. | <p>a. Enter the following command:</p> <pre><b>cifs terminate -t nn</b></pre> <p><i>nn</i> is a notification (in seconds) appropriate for your clients. After that period of time, proceed to Step 5.</p> <p>b. Wait for <i>nn</i> seconds and then go to Step 5.</p> |

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.

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

```
bye
```

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

11. Enter the following command:

```
update_flash
```

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

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

```
bye
```

13. 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. 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.

### After you finish

If desired, reenable automatic giveback.

## Updating system firmware using the disruptive method

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

### Before you begin

You must have obtained the system firmware from the N series support site on your Windows or UNIX client or your HTTP server.

### 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 N series support site.

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

```
cf enable
```

## Disk firmware updates

Disk firmware is bundled with the Data ONTAP system files and updated automatically during Data ONTAP upgrades. You can also obtain and update disk firmware manually; doing so is a prerequisite for major nondisruptive upgrades.

### How disk firmware is updated

When you upgrade Data ONTAP, disk firmware is updated automatically if the firmware on the disks is older than the firmware that is bundled with the Data ONTAP system files. You can also update disk firmware by downloading the most recent firmware package from the N series support site and installing the files.

Each storage system is shipped with an `/etc/disk_fw` directory that contains the latest firmware revisions. Disk firmware is updated automatically when one of the following is true:

- You add new disks or a disk shelf.  
Disk firmware updates are applied from the `/etc/disk_fw` directory.  
**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.

Disk firmware updates can be added to the `/etc/disk_fw` 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.
- After obtaining a disk firmware update package  
You might be directed to download a disk firmware update from the N series support site in the event that you encounter problems with certain disk types or you receive a notice from IBM.

You must download and install the latest disk firmware before upgrading Data ONTAP using the NDU method.

- When you hot-add a SAS shelf

Automatic background disk firmware updates are enabled by the `raid.background_disk_fw_update.enable` option, which is set to `on` by default. Do not change the default value unless you are directed to do so by technical support.

Automatic background disk firmware updates are overridden when the `disk_fw_update` command is issued. This command makes disks inaccessible for up to two minutes. It is not recommended that you use this command unless you are directed to do so by technical support.

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

By default, disk firmware updates take place automatically in the background, thus ensuring the continuity of storage system services. Nonetheless, you must download and install the latest disk firmware before performing an NDU.

In Data ONTAP 8.0.2 and later releases, automatic background disk firmware updates are available for non-mirrored RAID4 aggregates, in addition to all other RAID types. If your system includes non-mirrored RAID4 aggregates, it is no longer necessary to perform a disruptive disk firmware update before upgrading Data ONTAP.

## Detecting outdated disk firmware using AutoSupport

AutoSupport messages include information about disk firmware installed on your storage system. The 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 Installed Systems service to monitor disk firmware versions, your storage system must meet the following requirements:

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

### About this task

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

## Steps

1. Use a web browser to go to the N series support site at [www.ibm.com/storage/support/nseries/](http://www.ibm.com/storage/support/nseries/).
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 a firmware update is available for your storage system, you see a message with a link to a Firmware Analysis page. If the Firmware Analysis page contains a notice that newer disk firmware is available for your system, a disk firmware update takes 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.

# Updating disk shelf firmware

Disk shelf firmware is bundled with the Data ONTAP system files and updated automatically during Data ONTAP upgrades. You can also obtain and update disk shelf firmware manually; doing so is a prerequisite for major nondisruptive upgrades.

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

## Related concepts

[How disk shelf firmware is updated](#) on page 75

## 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 N series support site and installing the files.

The module (AT series, ESH series, or SAS IOM series) in a disk shelf provides for the interconnect of the disks to the host bus adapter interface, including signal integrity when disks are swapped. 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 can also be internal components in N3300, N3400, and N3600 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 is 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 N series support site if you plan to perform a nondisruptive upgrade of Data ONTAP software, or if you receive a notice from IBM.
- When you hot-add a SAS shelf

Data ONTAP scans the `/etc/shelf_fw` directory for new firmware every two minutes. 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.
- N series Health Trigger 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.

## Disk shelf firmware requirements for Data ONTAP nondisruptive upgrades

When you plan a Data ONTAP NDU, you must ensure that any disk shelves attached to your system meet NDU requirements for service availability *before* upgrading Data ONTAP.

The availability of storage system services during a disk shelf firmware update depends on the type of shelf modules your system uses. Updates to disk shelf firmware (firmware for modules on disk shelves) must be nondisruptive to ensure a successful Data ONTAP NDU.

The following table summarizes service availability during disk shelf firmware upgrades.

| If your disk shelf type is... | Shelf firmware can be updated nondisruptively <i>apart from</i> Data ONTAP NDU? | Shelf firmware can be updated nondisruptively <i>during</i> Data ONTAP NDU? |
|-------------------------------|---------------------------------------------------------------------------------|-----------------------------------------------------------------------------|
| ESH or external SAS           | Yes                                                                             | Yes, but updating before Data ONTAP NDU is a best practice                  |

| If your disk shelf type is... | Shelf firmware can be updated nondisruptively <i>apart from</i> Data ONTAP NDU? | Shelf firmware can be updated nondisruptively <i>during</i> Data ONTAP NDU?     |
|-------------------------------|---------------------------------------------------------------------------------|---------------------------------------------------------------------------------|
| Internal SAS                  | Yes, if requirements are met                                                    | Yes, if requirements are met; updating before Data ONTAP NDU is a best practice |
| AT-FCX                        | Yes, if requirements are met                                                    | No, must be updated before Data ONTAP NDU                                       |
| AT-FC or AT-FC2               | No, system downtime is required                                                 | No, system downtime is required                                                 |

- ESH-based and external SAS-based shelves  
ESH and SAS firmware updates are nondisruptive and do not result in I/O disruption during a Data ONTAP NDU. Nonetheless, it is a best practice for ESH and SAS modules to be running the latest firmware before an NDU. You should verify the current firmware version and update it if necessary before the Data ONTAP NDU.
- SAS internal shelves in N3300, N3400, or N3600 systems  
SAS firmware for modules in N3300, N3400, or N3600 systems can be updated nondisruptively separately from a Data ONTAP NDU when firmware version 0500 or later is running. If firmware version 0400 or earlier is running, you must schedule system downtime to update AT-FCX firmware. You must therefore verify the current firmware version to determine whether a Data ONTAP NDU is possible.
- AT-FCX shelves  
AT-FCX firmware can be updated nondisruptively separately from a Data ONTAP NDU when the following conditions are met:
  - Data ONTAP 7.3.2 or later is running.
  - Firmware version 37 or later is running.
  - Multipath Storage is implemented.

If these conditions are not met, you must schedule system downtime to update AT-FCX firmware.

However, after these conditions are met, you must verify the current firmware version and update it if necessary before the Data ONTAP NDU. That is because new AT-FCX firmware is bundled with Data ONTAP software images. When new AT-FCX firmware is installed automatically during a Data ONTAP upgrade, I/O is disrupted on each shelf when the modules are rebooted. Multipath Storage can be implemented in HA pairs or stand-alone systems. AT-FCX firmware can be upgraded nondisruptively in either configuration.
- AT-FC and AT-FC2 shelves  
Firmware for AT-FC and AT-FC2 modules cannot be updated nondisruptively; you must schedule system downtime. You should therefore verify the current firmware version to determine if you want system downtime to occur during the Data ONTAP upgrade or before.

## Detecting outdated disk shelf firmware

If you want to perform a nondisruptive upgrade of Data ONTAP software, 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.

### Steps

1. Go to the disk shelf firmware information on the N series support site and determine the most recent firmware version for your shelves.
2. At the storage system command line, enter the following command:  
`sysconfig -v`
3. 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
```

If the disk shelf firmware version in the command output is earlier than the most recent version on the N series support site, you must update your disk shelf firmware manually.

## Updating disk shelf firmware manually

You must run the `storage download shelf` command after downloading new disk shelf firmware from the N series support 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 N3400 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 prevents 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 N series support site.
2. Contact IBM support for instructions to extract your firmware files to the `/etc/shelf_fw` directory in the root volume of your storage system.
3. 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"> <li>• ESH-based disk shelves</li> <li>• SAS-based disk shelves</li> <li>• AT-FCX-based disk shelves running firmware version 37 or higher</li> <li>• N3400 internal shelves</li> </ul> | Go to the next step.                                                                                                                                                                                           |
| <ul style="list-style-type: none"> <li>• AT-based disk shelves</li> <li>• AT-FCX-based disk shelves running firmware version 36 or lower</li> </ul>                                                                       | Enter the following command:<br><b>cifs terminate -t <i>nn</i></b><br>where <i>nn</i> is a notification period (in minutes) appropriate for your clients. After that period of time, proceed to the next step. |

4. Enter the following command at the storage system console to access the advanced administrative commands:

```
priv set advanced
```

5. Enter the following command to update disk shelf firmware on all disk shelves in your system:

```
storage download shelf
```

If necessary, you can update the firmware only on shelves attached to a specific adapter. However, all disk shelves attached to a storage system should be running the same firmware version.

6. Enter the following command to verify the new disk shelf firmware:

```
sysconfig -v
```

7. Enter the following command to return to the standard administrative console prompt:

```
priv set admin
```

8. If you terminated CIFS before updating shelf firmware, reenable it by entering the following command:

```
cifs restart
```

## Updating ACP firmware

If your disk shelves include ACP functionality, ACP firmware can be updated automatically during Data ONTAP upgrades or manually after downloading new ACP processor firmware from the N series support site.

### Before you begin

- ACP interfaces must be cabled properly.
- 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 N series support site and installing the files using the `storage download acp` command.

Installing ACP firmware can take several minutes, but it does not disrupt client access. However, normal ACP recovery capabilities are not available while the firmware upgrade is in progress.

### Steps

1. Find and download the most recent ACP firmware on the N series support site.
2. Contact IBM support for instructions 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
```

4. Enter the following command to verify the new ACP firmware:

```
storage show acp
```

You 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<br>(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 is completed, you see output similar to the following:

| 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 |

## Service Processor firmware updates

Service Processor (SP) is a remote management device that is included in N6200 series and N7x50T series 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*.

### 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

#### Step

1. Go to Firmware Instructions for the Service Processor at the N series support site and follow the instructions provided.

### 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

[www.ibm.com/storage/support/nseries/](http://www.ibm.com/storage/support/nseries/)

## 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

### Step

1. Go to Firmware Instructions for the Service Processor at the N series support site and follow the instructions provided.

### Related information

[www.ibm.com/storage/support/nseries/](http://www.ibm.com/storage/support/nseries/)

## 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*.

## 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 |

- Complete the steps as directed in the following table based on your RLM firmware version.

| If the firmware version is... | Then...                                                                                                                                                         |
|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Earlier than 3.1</b>       | Complete Steps 3 through 7 to upgrade to the latest 3.1.x version.<br>If you want to update to version 4.0 or later, you must also complete Steps 8 through 13. |
| <b>3.1.x</b>                  | Complete Steps 8 through 13.                                                                                                                                    |
| <b>4.0 or later</b>           | Complete Steps 3 through 7.                                                                                                                                     |

- Go to Firmware Instructions for the Remote LAN Module at the N series support site.
- Click the `RLM_FM.zip` link to download the file from the N series support 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 N series support site is the same as the version running on your RLM, it is not necessary to update RLM firmware at this time.

- Enter the following command at the storage system prompt:

```
software update http://Web_server/RLM_FW.zip -f
```

- 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.

- 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.  |

- 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 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.

12. 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.

13. Enter the following command to return to the standard administrative console prompt:

```
priv set admin
```

## 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

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 |

- Complete the steps as directed in the following table based on your RLM firmware version.

| If the firmware version is... | Then...                                                                                                                                                         |
|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Earlier than 3.1</b>       | Complete Steps 3 through 7 to upgrade to the latest 3.1.x version.<br>If you want to update to version 4.0 or later, you must also complete Steps 8 through 13. |
| <b>3.1.x</b>                  | Complete Steps 8 through 11.                                                                                                                                    |
| <b>4.0 or later</b>           | Complete Steps 3 through 7.                                                                                                                                     |

- Go to Firmware Instructions for the Remote LAN Module at the N series support site.
- Click the `RLM_FM.tar.gz` link to download the file from the N series support 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 N series support 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.

## 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 more information about the Event Management System (EMS), see the `ems(1)` man page.

## 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 83

## 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 85

## 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 disruptive 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*.

### Related concepts

[Installing Data ONTAP software images](#) on page 43



## 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 disruptive 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:

| <b>If CIFS...</b>          | <b>Then...</b>                                                                                                                                                                                                                              |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Is not in use in system A. | Go to Step 3.                                                                                                                                                                                                                               |
| Is in use in system A.     | Enter the following command:<br><br><b>cifs terminate -t nn</b><br><br><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 disruptive method

The disruptive 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
```

## Flash Cache firmware updates

Firmware for Flash Cache devices is included with distribution files for upgrades to Data ONTAP 8.0.1 and later releases. If the running firmware is older than the firmware that is bundled with the Data ONTAP system files, it is updated automatically.

Firmware updates are not available for the original 16-GB PAM devices. Automatic updates occur only to Flash Cache devices.

If you are upgrading Data ONTAP nondisruptively (NDU), Flash Cache firmware is updated nondisruptively. This is because the reboot required for Flash Cache firmware upgrades occurs before the final reboot of the `cf giveback` process. Consequently, if your system includes Flash Cache devices, you might see multiple reboots during a Data ONTAP NDU; this is expected behavior.

For information about what Flash Cache and PAM are and how they work, see the *Data ONTAP 7-Mode System Administration Guide*.



## Reverting to an earlier 7-Mode release family

---

Transitioning a storage system to a Data ONTAP 7-Mode release in an earlier family is referred to as a *reversion*. Reverting requires preparation, using `revert_to` command, and completing post-reversion procedures.

The `revert_to` command modifies Data ONTAP on-disk structures to be compatible with the earlier target release and ensures that the system is prepared for the reversion.

**Attention:** *Do not* attempt to revert Data ONTAP by simply downloading and booting (or netbooting) a release in an earlier release family. If you do, you cannot boot the earlier target release. You must use the `revert_to` command for the reversion process.

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

## When to revert and when to call technical support

You can revert without assistance when reverting new or test systems, but you should call technical support if you encounter problems during or after upgrading.

You can revert to an earlier release family without assistance from technical support only in the following scenarios:

- You upgraded to a new release on a test system and you want to return to the original release when testing is completed.
- You are configuring a new storage system — running a later release of Data ONTAP and not yet in production — in an environment in which you have standardized on an earlier Data ONTAP release.

*Do not* attempt to revert Data ONTAP without assistance in the following circumstances:

- The upgrade process fails and cannot finish.
- The upgrade process finishes, but the system is unusable in a production environment.
- The upgrade process finishes and the system goes into production, but you are not satisfied with its behavior.

In these circumstances, contact technical support immediately.

## 7-Mode reversion checklist

To ensure a successful reversion, you must check several things before, during, and after the reversion.

### Preparing to revert

Preparatory steps are complete when all of the following conditions are true:

- General reversion requirements have been satisfied, including:
  - Software and hardware support in the target release is confirmed.
  - System status requirements are addressed.
  - Standard system services and processes that could interfere with reversion are not running.
- For production systems being reverted, potential issues in your environment have been identified and preparatory tasks completed as appropriate.
- The target Data ONTAP image has been obtained from The N series support site and is available to systems being reverted.

### Performing the reversion

Software reversion steps are complete when all of the following conditions are true for each partner in an HA pair:

- Any remaining conditions identified by the `revert_to` command have been addressed.
- The `revert_to` command has completed and each system has booted the target release.
- The correct SP firmware is loaded and running on supported platforms.

### After reverting

Post-reversion steps are complete when all of the following conditions are true:

- HA relationship is restored between partner nodes.
- Services and protocols are functioning as expected.
- For production systems being reverted, potential issues in your environment have been identified and post-reversion tasks completed as appropriate.

## General reversion requirements

You must satisfy these requirements before you revert to a previous Data ONTAP version. If the `revert_to` command encounters one of these conditions, it halts until the condition is addressed.

### Target release requirements

- You cannot revert to a release earlier than Data ONTAP 7.3.



- You must disable any 8.0 release family features before reverting.
- You cannot revert a Data ONTAP 8.0 system to an earlier release family if it contains 64-bit aggregates.  
The Data ONTAP 7.3 release family does not support 64-bit aggregates. You must remove any 64-bit aggregates before reverting to a 7.3.x release. You should contact technical support before attempting to revert a production system that contains 64-bit aggregates.
- 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.  
If you upgraded Data ONTAP for new hardware support, you must disconnect the new hardware and reconfigure your system before reverting.
- You must verify that all components of your configuration are compatible with the target Data ONTAP reversion release by consulting the the compatibility and configuration information about FCP and iSCSI products.  
See the appropriate matrix at the N series Service and Support website at [www.ibm.com/storage/support/nseries/](http://www.ibm.com/storage/support/nseries/).

### System status requirements

- All disks must be online.
- All volumes and aggregates 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.
- All volumes and aggregates must be free of file system errors and bad blocks.
- Compressed volumes must be decompressed for reversion.
- Any unsaved core must be recovered or released
- Snapshot copies made on Data ONTAP release families later than the target release cannot be present on the source system. Before reverting to an earlier release family, you must delete them.
- All SnapMirror relationships must be broken.
- Metadata for deduplicated volumes must be reverted or removed using the `sis revert_to` command.
- All LUNs in the system must be owned by the default vfiler `vfiler0`
- You cannot revert if an upgrade is in progress. You must complete the upgrade before reverting.
- You cannot revert if the background quota upgrade is still in process from a previous Data ONTAP upgrade.

### Operational requirements

The `revert_to` command halts with an error message if any of these conditions are encountered. You can re-enter the `revert_to` after addressing them.

- The following services and protocols cannot be running during a Data ONTAP reversion:
  - High availability (HA) takeover and giveback
  - NFS

- CIFS
- FCP
- SnapMirror
- SnapVault
- The following jobs cannot be running during a Data ONTAP reversion:
  - Dump or restore
  - RAID scrubs
  - RAID optimized reconstructions
  - RAID assimilation
  - RAID disk sanitization
  - waflliron
  - Inode file upgrade
  - Disk maintenance center testing
  - Disk failure processing

## Requirements for reverting configured systems

If you configured Data ONTAP features on a new storage system after initial system setup, or if you upgraded a system and modified the configuration, you must satisfy certain requirements before you revert the system to a previous Data ONTAP version, in addition to the general reversion requirements.

### Issues to address *before* reverting

You must evaluate your needs for the following Data ONTAP capabilities in the target system and if necessary, prepare your system before reverting:

- storage capacity  
Your storage system must conform to the maximum capacity limitations of the earlier release.
- space guarantees  
Space guarantees do not persist through reversions to earlier releases.

If you have implemented any of the following Data ONTAP capabilities, you must evaluate the requirements and if necessary, modify your configuration before reverting.

- SnapMirror
- interface group configuration in the `/etc/rc` file
- VLANs
- deduplication
- SSDs
- Brocade switches in fabric-attached MetroClusters
- FlexClone technology

### Issues to address *after* reverting

You must evaluate your needs for the following Data ONTAP capabilities in the target system and if necessary, adjust your configuration after reverting:

- deduplication
- NDMP
- volumes with rewritten FSIDs
- TOE
- in-order frame delivery on FC switches

## 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`

## Preparing to revert Data ONTAP 7-Mode

Before reverting to an earlier Data ONTAP release family, you must verify reversion requirements, resolve any reversion issues, and obtain the Data ONTAP software image for the target release.

Be sure to check the *Release Notes* for this Data ONTAP source release for any updates to reversion notices and procedures.

## Commands for addressing reversion requirements

You can view the status of system conditions and operations that affect Data ONTAP reversions and take appropriate action before issuing the `revert_to` command.

### System status

| To verify that...                    | Use this command to check status... | Address the requirement before reverting by...                                                                                                                                                                                                                                                                                                                                      |
|--------------------------------------|-------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| No 64-bit aggregates are present     | <code>aggr status</code>            | <p>Migrating the data to 32-bit aggregates using Qtree SnapMirror and destroy the 64-bit aggregates.</p> <p>For more information about Qtree SnapMirror, see the <i>Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide</i>.</p> <p><b>Attention:</b> Contact technical support before attempting to revert a production system that contains 64-bit aggregates.</p> |
| No disks are offline                 | <code>aggr status -f</code>         | Bringing them online or replacing them.                                                                                                                                                                                                                                                                                                                                             |
| No volumes are offline or restricted | <code>vol status</code>             | <p>Using one of the following commands:</p> <ul style="list-style-type: none"> <li>• <code>vol online</code></li> <li>• <code>vol destroy</code></li> </ul>                                                                                                                                                                                                                         |
| No volumes are compressed            | <code>vol compress status</code>    | Using the <code>vol decompress revert</code> command.                                                                                                                                                                                                                                                                                                                               |

| To verify that...                                       | Use this command to check status...  | Address the requirement before reverting by...                                                                                                                                                                                                                                                               |
|---------------------------------------------------------|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Deduplicated volumes are prepared for reversion         | <code>sis revert_to 7.3</code>       | Running one of the following commands: <ul style="list-style-type: none"> <li>• <code>sis revert_to 7.3</code>, to revert deduplication metafiles</li> <li>• <code>sis revert_to 7.3 -delete</code>, to delete deduplication metafiles</li> </ul> See the <code>sis(1)</code> man page for more information. |
| No volumes are marked <code>waf1 inconsistent</code>    | <code>vol status</code>              | Contact technical support immediately.                                                                                                                                                                                                                                                                       |
| No aggregates are offline or restricted                 | <code>aggr status</code>             | Using one of the following commands: <ul style="list-style-type: none"> <li>• <code>aggr online</code></li> <li>• <code>aggr destroy</code></li> </ul>                                                                                                                                                       |
| No aggregates are marked <code>waf1 inconsistent</code> | <code>aggr status</code>             | Contact technical support immediately.                                                                                                                                                                                                                                                                       |
| No unsaved cores are present                            | <code>savecore -i</code>             | Recover them or release them with the <code>savecore</code> command                                                                                                                                                                                                                                          |
| No SnapShot copies are present                          | <code>snap list</code>               | Deleting them with the <code>snap delete</code> command                                                                                                                                                                                                                                                      |
| No SnapMirror relationships are present                 | <code>snapmirror destinations</code> | Breaking them with the <code>snapmirror break</code> command                                                                                                                                                                                                                                                 |
| Any previous upgrade has completed.                     | n.a.                                 | Waiting at least 10 minutes after an upgrade before beginning a reversion.                                                                                                                                                                                                                                   |
| No background quota upgrade is in process               | <code>quota status</code>            | Disabling quotas or allowing the quota upgrade to complete                                                                                                                                                                                                                                                   |

## Services and protocols

| To ensure that the following services are not running... | Use this command to check status... | Use this command to halt the operation manually... |
|----------------------------------------------------------|-------------------------------------|----------------------------------------------------|
| High availability (HA) takeover and giveback             | <code>cf status</code>              | <code>cf disable</code>                            |
| NFS                                                      | <code>nfs status</code>             | <code>nfs stop</code>                              |
| CIFS                                                     | <code>cifs status</code>            | <code>cifs terminate</code>                        |
| FCP                                                      | <code>fcg status</code>             | <code>fcg stop</code>                              |
| iSCSI                                                    | <code>iscsi status</code>           | <code>iscsi stop</code>                            |
| SnapMirror                                               | <code>snapmirror status</code>      | <code>snapmirror off</code>                        |
| SnapVault                                                | <code>snapvault status</code>       | <code>snapvault stop</code>                        |

## Jobs

If any of these jobs are running, you can halt them manually or you can wait until the operation finishes.

| To ensure that the following jobs are not running... | Use this command to check status...                                                                    | Use this command to halt the operation manually... |
|------------------------------------------------------|--------------------------------------------------------------------------------------------------------|----------------------------------------------------|
| Dump or restore                                      | <code>backup status</code>                                                                             | <code>backup terminate</code>                      |
| RAID scrubs                                          | <code>aggr scrub status</code>                                                                         | <code>aggr scrub stop</code>                       |
| RAID optimized reconstructions                       | <code>aggr status</code>                                                                               | Allow the operation to finish.                     |
| RAID disk sanitization                               | <code>disk sanitize status</code>                                                                      | <code>disk sanitize abort</code>                   |
| wafiron                                              | wafiron should only be run under direction from technical support; consult with them before reverting. |                                                    |
| Inode file upgrade                                   | <code>*waf1 scan status</code> (this is an advanced command)                                           | Allow the scan to finish.                          |
| Disk maintenance center testing                      | <code>disk maint status</code>                                                                         | <code>disk maint abort</code>                      |
| Disk failure processing                              | <code>disk show -v</code> or <code>storage show disk -a</code>                                         | Identify and remove any failed disks.              |

## Related concepts

*When to revert and when to call technical support* on page 95

## Preparing to revert configured systems

If you are reverting a system that you have configured to serve data to clients in your environment, you must ensure that certain configurations are prepared for the reversion.

### Requirements for reverting to a Data ONTAP release with a lower maximum capacity

When you revert to an earlier Data ONTAP release, your storage system must conform to the maximum capacity limitations of the earlier release.

If you upgraded your system to a release that supports greater capacities and you configured storage to utilize the new capacities, you must reconfigure your system to the lower capacity limits before you revert. If you don't reconfigure in this way, the storage system does not boot up following the revert process until the excess capacity has been disconnected.

You can reduce the total capacity of your system by destroying aggregates or by moving aggregates to a different system. The system to which storage is relocated must meet the following requirements:

- It has spare capacity to accommodate the relocated storage.
- It is running the same or a later Data ONTAP release as the system where the disks are currently installed.
- It is running a Data ONTAP release that supports the relocated disks.

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

For more information about maximum capacity limits for a given Data ONTAP release, see the appropriate hardware and service guide entries for that release.

### Considerations for reverting systems with space guarantees enabled

Space guarantees do not persist through reversions to earlier Data ONTAP software versions. Before reverting a system with space guarantees enabled, you should review your configuration to ensure that space guarantees behave as expected.

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.

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*.

## 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.

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.

- |                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Volume SnapMirror</b> | 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.                                    |
| <b>Qtree SnapMirror</b>  | 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.



## 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
```

## Considerations for reverting a system with deduplicated volumes

In addition to reverting deduplication metafiles with the `sis revert_to` command, which is required by the `revert_to` command, you must also verify volume size and SnapMirror requirements.

Before reverting a system with deduplicated volumes, ensure that you have satisfied the following requirements:

- You must ensure adequate free space in the deduplicated volumes.  
Running the `sis revert_to` command creates downgraded copies of existing deduplication metafiles. Before running the command, you must therefore ensure that 4 to 6 percent of the logical data size in the volume is available for the new files.
- If the release you are reverting to does not support your current volume size, you must decrease the size of any deduplicated volume to the limit supported in the target release.

- For deduplication volumes that are replicated using SnapMirror, the destination storage system should support deduplication.
- Prepare the duplicated volumes for reversion by running the `sis revert_to` command. The general `revert_to` command requires that deduplication metafiles either be downgraded or removed. Running the command without options retains the old files and creates new ones, running it with the `-delete` option deletes the existing metafiles.

To continue using deduplication when the Data ONTAP reversion is complete, you must enable the NearStore license. If you deleted the deduplication metadata, you must also rebuild the deduplication metadata.

For more information about reverting systems with deduplicated volumes, see the `sis(1)` man page.

### Related tasks

[Using deduplication on a reverted system](#) on page 113

### Related references

[Commands for addressing reversion requirements](#) on page 100

## 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 functionality 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 functionality 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 functionality 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 functionality license and the storage systems that support deduplication, see the *Data ONTAP 7-Mode Storage Management Guide*.

## 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*.

## 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 N series support website.

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.

## 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 to an earlier release family removes IBM customer contact information

The IBM customer contact information is removed from the registry when a storage system running Data ONTAP 8.0 is reverted to any 7.3.x or 7.2.x release. To work around this problem, you must record your contact information before reverting and reenter it manually after reverting.

### About this task

For more information about IBM contact and location information for your storage system, see the *Data ONTAP 7-Mode System Administration Guide*.

### Steps

1. On the system running Data ONTAP 8.0, enter the following command to list customer contact information:  

```
options contact
```
2. Revert the system as directed by technical support.
3. When the system has booted the earlier release, enter the contact information.

### Related concepts

[Updating IBM customer contact information](#) on page 67

## 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.

## 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.

## Staging the target Data ONTAP image

You must obtain the software image for the target Data ONTAP reversion or downgrade release and make it accessible to the storage system.

### About this task

You can use your preferred method -- HTTP, UNIX client, or Windows client access -- to make the reversion target image available to the storage system.

### Step

1. Copy the target software image (for example, `80_setup_i.tgz`) from the N series support site or another storage system to the HTTP server or client system you use to stage software images.

### Related concepts

[Installing Data ONTAP software images](#) on page 43

## Performing the 7-Mode reversion process

To revert to an earlier 7-Mode Data ONTAP release, you must halt certain processes, install the target image, and enter the `revert_to` command. If your system includes a Service Processor, you must verify its firmware version.

## Reverting Data ONTAP

To revert Data ONTAP, you must install the target release on your system and run the `revert_to` command.

### Before you begin

You must obtain the target Data ONTAP image and stage it either on a web server that is accessible to your storage system or in the `/etc/software` directory on the storage system.

You should verify that protocol, system services, and RAID operations are not running before proceeding with this task. If any are running, the `revert_to` command halts and prompts you to correct the condition before proceeding.

**About this task**

You must revert the partner system in an HA pair before booting the systems into the earlier target release.

**Steps**

1. Enter the following command to install the target Data ONTAP image on your system and commit it to the boot device:

```
software update url/file
```

Supply the URL of a web server where you staged the image.

When prompted, confirm that you want to downgrade the file system and perform a revert.

You see output similar to the following:

```
software: You can cancel this operation by hitting Ctrl-C in the next 6
seconds.
software: Depending on system load, it may take many minutes
software: to complete this operation. Until it finishes, you will
software: not be able to use the console.
software: copying to 735_setup_q.zip
software: 100% file read from location.
software: /etc/software/735_setup_q.zip has been copied.
software: installing software, this could take a few minutes...
software: Data ONTAP(R) Package Manager Verifier 1
software: Validating metadata entries in /etc/boot/NPM_METADATA.txt
software: Downgrade WAFL from version 21057 to 19744 (Y/N)? Y
software: The release that you are installing will Downgrade the file
system
causing the need to use the revert_to command BEFORE a reboot. Do you
wish to continue?
(Y/N) Y
software: Checking sha1 checksum of file checksum file:
 /etc/boot/NPM_FCSUMx86-64.shal.asc
software: Checking sha1 file checksums in /etc/boot/NPM_FCSUM-
x86-64.shal.asc
software: installation of 735_setup_q.zip completed.
Thu Oct 28 17:45:14 GMT [filer: cmds.software.installDone:info]:
Software:
Installation of 735_setup_q.zip was completed.
...
download: Downloading boot device (Service Area)
...
software: Due to WAFL downgrade, user should use "revert_to" instead of
"reboot"
Please type "revert_to" for the changes to take effect.
Thu Oct 28 17:45:19 GMT [filer: download.requestDone:notice]: Operator
requested
download completed
```

2. Confirm that HA takeover/giveback is disabled by entering the following command:

```
cf status
```

If the HA relationship is still enabled, enter the `cf disable` command before proceeding.

3. Enter one of the following commands, depending on your environment:

| If deduplication is... | Then enter...                        |
|------------------------|--------------------------------------|
| <b>Deployed</b>        | <b><code>revert_to -f 7.3</code></b> |
| <b>Not deployed</b>    | <b><code>revert_to 7.3</code></b>    |

Confirm that you want to proceed when prompted.

If you have not satisfied the prerequisites, the reversion process halts and you are prompted to address the problem. When you have done so, you can restart the `revert_to` command.

You see output similar to the following:

```
You are about to revert the system to work with Data ONTAP 7.3
The system will be halted immediately after the conversion
process completes. Make sure that you have installed Data ONTAP
7.3 onto the boot device, or you will have to run "revert_to" again.

Are you sure you want to proceed? [yes/no]? yes
Mon Nov 15 17:50:24 GMT [filer: revertTo.start:notice]: Starting revert
to
7.3.
...
Reboot the system with Data ONTAP 7.3[.x].
Thu Oct 28 17:50:50 GMT [filer: revertTo.complete:notice]: Revert to
7.3[.x]
was completed.
Setting boot image to 7G.
Clearing next boot image
...
RAID revert complete. You can reboot the system after partner has been
reverted.
System halting...
```

When the system has halted, the `LOADER` prompt is displayed.

4. If the system is a partner in an HA pair, repeat Steps 1 and 2 on the partner system.
5. At the `LOADER` prompt of each partner node, enter the following command to reset the system BIOS:

**`bye`**

Each node autoboots normally on the target Data ONTAP 7.3.x release.

After booting for the first time on a Data ONTAP 7.3.x release after reverting the system from Data ONTAP 8.0.x, you might encounter a large number of errors (repeated once for each disk in the system) of the following form:

```
[filer: diskown.RescanMessageFailed:warning]: Could not send rescan
message to filer. Please type disk show on the console of filer for it
to scan the newly inserted disks.
```

This message will be repeated once for each disk owned by the system. This is a known issue that can safely be ignored; the error will only occur on the first boot after reverting.

### After you finish

If your system includes a Service Processor (SP), ensure that its firmware is up to date. Otherwise, proceed to post-reversion tasks.

## Updating SP firmware

If your storage system includes a Service Processor (SP), you must verify that it is running the correct firmware version and update it if it is not.

### Before you begin

The reversion or downgrade process should be complete and the storage system should be running the target release.

### About this task

Data ONTAP software images include firmware for SP modules. If the firmware version on your SP module is outdated, you must update it before returning the reverted or downgraded system to production.

### Steps

1. Go to the system firmware information on the N series support site and determine the most recent firmware version for your SP module.
2. Enter the following command at the storage system CLI to determine the SP firmware version:

```
sp status
```

You see output similar to the following:

```
Service Processor Status: Online
Firmware Version: 1.2
...
```

If the SP firmware version in the command output is earlier than the most recent version on the N series support site, you must update your disk shelf firmware manually.

3. Click the `SP_FW.zip` link to download the file from the N series support site to your HTTP server.
4. At the storage system prompt, enter the following command:
 

```
software update http://Web_server/SP_FW.zip -f
```
5. When the `software update` command is finished, enter the following command at the storage system prompt:



```
sp update
```

6. When the system prompts you to update SP, enter **y** to continue.

SP is updated and you are prompted to reboot SP. Wait approximately 60 seconds to allow SP to reboot.

7. Verify that the SP firmware has been updated by entering the following command:

```
sp update
```

8. If the system is a partner in an HA pair, repeat Steps 4 through 7 on the partner system.

### Result

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.

## Completing post-reversion tasks

After reverting to an earlier Data ONTAP release family, you might need to perform additional tasks.

You should verify that any services you halted manually restarted after the reversion. If not, you should restart them manually and verify that any clients have appropriate access to storage system services.

## Using deduplication on a reverted system

If your storage system is licensed for deduplication and you revert the system to a release in the Data ONTAP 7.3 family, you must add a NearStore license to reenable deduplication. If you deleted deduplication metafiles before reverting, you must rebuild them to continue running deduplication.

### Steps

1. To determine if NearStore is licensed on the reverted system, enter the following command:

```
license
```

If the `nearstore_option` entry displays `not licensed`, you must add the license code using the `license add` command.

If you do not have a NearStore license, contact your sales or support representative to obtain one.

2. If you deleted the deduplication metafiles before reverting, run the following command on every deduplicated volume to rebuild them:

```
sis start -s path
```

This process can take several days, depending on the size of the logical data in the volume.

For more information about rebuilding deduplication metadata, see the `sis(1)` man page.

## Reenabling NDMP on a reverted system

The reversion process removes the registry entry that enables NDMP after a storage system reboots. If you want to continue using NDMP after the reversion, you must reenoble it manually.

### Steps

1. To reenoble NDMP immediately, enter the following command:

```
ndmpd on
```

2. To enable NDMP at each system reboot, enter the command `ndmpd on` in the `/etc/rc` file.

## 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 reenoble 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.

## Reinstatement of in-order frame delivery after reversion

If out-of-order frame delivery is enabled and you revert to a previous Data ONTAP release, the out-of-order frame delivery functionality is disabled. You must manually enable out-of-order frame delivery after reverting to a Data ONTAP release that supports this functionality.

If you are reverting to a Data ONTAP release that does not support out-of-order frame delivery, you must manually enable the in-order frame delivery options and port-based policy on FC switches.

For more information about enabling out-of-order frame delivery, see the *Data ONTAP 7-Mode Data Protection Online Backup and Recovery Guide* and your FC switch documentation.



## Downgrading to an earlier release in the same 7-Mode release family

---

Use the downgrade process when you need to transition your storage system to an earlier Data ONTAP 7-Mode release in the same release family. Doing so requires preparation, downloading and booting the earlier release, and completing post-downgrade procedures.

Transitioning a storage system to an earlier Data ONTAP release in the same release family is referred to as a *downgrade*. Downgrading does not require modifications to Data ONTAP on-disk structures, you simply need to obtain and boot the target release after verifying requirements and compatibility.

### When to downgrade and when to call technical support

You can downgrade without assistance when downgrading new or test systems, but you should call technical support if you encounter problems during or after upgrading.

You can downgrade to an earlier release family without assistance from technical support only in the following scenarios:

- You upgraded to a new release on a test system and you want to return to the original release when testing is completed.
- You are configuring a new storage system — running a later release of Data ONTAP and not yet in production — in an environment in which you have standardized on an earlier Data ONTAP release.

*Do not* attempt to downgrade Data ONTAP without assistance in the following circumstances:

- The upgrade process fails and cannot finish.
- The upgrade process finishes, but the system is unusable in a production environment.
- The upgrade process finishes and the system goes into production, but you are not satisfied with its behavior.

In these circumstances, contact technical support immediately.

### 7-Mode downgrade checklist

To ensure a successful downgrade, you must check several things before, during, and after the reversion.

#### Preparing to revert

Preparatory steps are complete when all of the following conditions are true:

- General downgrade requirements have been satisfied, including:
  - Software and hardware support in the target release is confirmed.
  - System status requirements are addressed.
  - Standard system services and processes that could interfere with downgrade are not running.
- For systems being downgraded to the Data ONTAP 8.0 release, potential issues in your environment have been identified and preparatory tasks completed as appropriate.
- The target Data ONTAP image has been obtained from The N series support site and is available to systems being reverted.

### **Performing the downgrade**

Software downgrade steps are complete when all of the following conditions are true for each partner in an HA pair:

- The target release has been downloaded to the boot device.
- Each system has booted the target release.
- The correct SP firmware is loaded and running on supported platforms.

### **After downgrading**

Post-reversion steps are complete when all of the following conditions are true:

- HA relationship is restored between partner nodes.
- Services and protocols are functioning as expected.

## **General downgrade requirements**

You must satisfy these requirements before you downgrade to a previous Data ONTAP version.

### **Target release requirements**

- You must disable any features not supported in the target release before downgrading.
- 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 downgrade to the earlier release. If you upgraded Data ONTAP for new hardware support, you must disconnect the new hardware and reconfigure your system before downgrading.
- You must verify that all components of your configuration are compatible with the target Data ONTAP downgrade release by consulting the the compatibility and configuration information about FCP and iSCSI products.  
See the appropriate matrix at the N series Service and Support website at [www.ibm.com/storage/support/nseries/](http://www.ibm.com/storage/support/nseries/).

### **System status requirements**

- All disks must be online.

- All volumes and aggregates must be online before downgrading.
- All volumes and aggregates must be free of file system errors and bad blocks.
- Any unsaved core must be recovered or released.
- All SnapMirror relationships must be broken.
- All LUNs in the system must be owned by the default vfiler vfiler0.
- You cannot downgrade if an upgrade is in progress. You must complete the upgrade before downgrading.
- You cannot downgrade if the background quota upgrade is still in process from a previous Data ONTAP upgrade.

### Operational requirements

The following jobs cannot be running during a Data ONTAP downgrade:

- Dump or restore
- RAID scrubs
- RAID optimized reconstructions
- RAID assimilation
- RAID disk sanitization
- wafliiron
- Inode file upgrade
- Disk maintenance center testing
- Disk failure processing

## Requirements when downgrading to Data ONTAP 8.0 7-Mode

There are additional requirements when downgrading from a later release to the Data ONTAP 8.0 release.

### 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.

#### About this task

Compression checkpoints are deleted during reversions.

#### 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 - * * * *
```

## 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 is 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 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.

## Preparing to downgrade Data ONTAP

Before downgrading to an earlier Data ONTAP release in the same release family, you must verify reversion requirements, resolve any downgrade issues, and obtain the Data ONTAP software image for the target release.

Plan to do the following:

- Read the *Release Notes* for this Data ONTAP source release.
- Verify that all components of your configuration are compatible with the target Data ONTAP downgrade release by consulting the the compatibility and configuration information about FCP and iSCSI products.

See the appropriate matrix at the N series Service and Support website at [www.ibm.com/storage/support/nseries/](http://www.ibm.com/storage/support/nseries/).



## Commands for addressing downgrade requirements

You must view the status of system conditions and operations that affect Data ONTAP downgrades and take appropriate action before performing the downgrade process.

### System status

| To verify that...                                       | Use this command to check status...                                        | Address the requirement before reverting by...                                                                                                     |
|---------------------------------------------------------|----------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| No disks are offline                                    | <code>aggr status -f</code>                                                | Bringing them online or replacing them.                                                                                                            |
| No volumes are offline or restricted                    | <code>vol status</code>                                                    | Using one of the following commands: <ul style="list-style-type: none"> <li><code>vol online</code></li> <li><code>vol destroy</code></li> </ul>   |
| No volumes are compressed                               | <code>vol compress status</code>                                           | Using the <code>vol decompress revert</code> command.                                                                                              |
| No volumes are marked <code>waf1 inconsistent</code>    | <code>vol status</code>                                                    | Contact technical support immediately.                                                                                                             |
| No aggregates are offline or restricted                 | <code>aggr status</code>                                                   | Using one of the following commands: <ul style="list-style-type: none"> <li><code>aggr online</code></li> <li><code>aggr destroy</code></li> </ul> |
| No aggregates are marked <code>waf1 inconsistent</code> | <code>aggr status</code>                                                   | Contact technical support immediately.                                                                                                             |
| No unsaved cores are present                            | <code>savecore -i</code>                                                   | Recover them or release them with the <code>savecore</code> command                                                                                |
| No SnapShot copies are present                          | <code>snap list</code>                                                     | Deleting them with the <code>snap delete</code> command                                                                                            |
| No SnapMirror relationships are present                 | <code>snapmirror destinations</code>                                       | Breaking them with the <code>snapmirror break</code> command                                                                                       |
| Any previous upgrade has completed.                     | Waiting at least 10 minutes after an upgrade before beginning a reversion. |                                                                                                                                                    |
| No background quota upgrade is in process               | <code>quota status</code>                                                  | Disabling quotas or allowing the quota upgrade to complete                                                                                         |

**Jobs**

If any of these jobs are running, you can halt them manually or you can wait until the operation finishes.

| To ensure that the following jobs are not running... | Use this command to check status...                                                                      | Use this command to halt the operation manually... |
|------------------------------------------------------|----------------------------------------------------------------------------------------------------------|----------------------------------------------------|
| Dump or restore                                      | <code>backup status</code>                                                                               | <code>backup terminate</code>                      |
| RAID scrubs                                          | <code>aggr scrub status</code>                                                                           | <code>aggr scrub stop</code>                       |
| RAID optimized reconstructions                       | <code>aggr status</code>                                                                                 | Allow the operation to finish.                     |
| RAID disk sanitization                               | <code>disk sanitize status</code>                                                                        | <code>disk sanitize abort</code>                   |
| waflliron                                            | waflliron should only be run under direction from technical support; consult with them before reverting. |                                                    |
| Inode file upgrade                                   | <code>*wafll scan status</code> (this is an advanced command)                                            | Allow the scan to finish.                          |
| Disk maintenance center testing                      | <code>disk maint status</code>                                                                           | <code>disk maint abort</code>                      |
| Disk failure processing                              | <code>disk show -v</code> or <code>storage show disk -a</code>                                           | Identify and remove any failed disks.              |

**Staging the target Data ONTAP image**

You must obtain the software image for the target Data ONTAP reversion or downgrade release and make it accessible to the storage system.

**About this task**

You can use your preferred method -- HTTP, UNIX client, or Windows client access -- to make the reversion target image available to the storage system.

**Step**

1. Copy the target software image (for example, `80_setup_i.tgz`) from the N series support site or another storage system to the HTTP server or client system you use to stage software images.

**Performing the 7-Mode downgrade process**

You can downgrade to an earlier 7-Mode Data ONTAP release using the nondisruptive or disruptive methods. If your system includes a Service Processor, you might need to update its firmware after the downgrade.

## Downgrading Data ONTAP using the nondisruptive method

You can downgrade HA pairs within a Data ONTAP release family while maintaining storage system availability. This nondisruptive downgrade 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 downgrade procedure, you must complete 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:<br><br><b>software update url/file</b><br><br>Then go to Step 4. |

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

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:<br><br><b>cifs terminate -t nn</b><br><br><i>nn</i> is a notification period (in minutes) 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:<br><br><b>cf giveback</b>                                                                                        |
| 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:<br><br><b>cf giveback</b> |

This command causes system A to reboot with the target Data ONTAP version and resume normal operation as a high-availability partner.

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 7 to update the partner storage system; in other words, bring down and update system B with partner A in takeover mode. |

| <b>If FCP or iSCSI...</b> | <b>Then...</b>                                                                                                                                                                                                                                                                                                   |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 7 to update the partner storage system; in other words, bring down and update system B with system A in takeover mode. |

## Downgrading Data ONTAP using the disruptive method

If you can schedule downtime to downgrade Data ONTAP, you can take HA pairs offline, then install the target release on your systems, download it to the boot device, and reboot.

### Before you begin

You must obtain the target Data ONTAP image and stage it on a web server that is accessible to your storage system.

You must verify that protocol, system services, and RAID operations are not running before proceeding with this task.

### About this task

You must downgrade the partner system in an HA pair before booting the systems into the earlier target release.

### Steps

1. If your systems are in an HA configuration, disable it 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:

| <b>If you...</b>                                                  | <b>Then...</b>                                                                                                               |
|-------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------|
| 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:<br><br><b>software update url/file</b><br><br>Then go to Step 4. |

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

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:

```
reboot
```

| If your systems are... | Then...                    |
|------------------------|----------------------------|
| Standalone             | The downgrade is complete. |
| In an HA pair          | Proceed to the next step.  |

5. While the HA configuration is disabled, repeat Step 2 through Step 4 at the console of system B.

**Attention:** Do not proceed to Step 6 until both systems in the HA configuration have been rebooted with the new version of Data ONTAP.

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

```
cf enable
```

**After you finish**

If your system includes a Service Processor (SP), ensure that its firmware is up to date. Otherwise, proceed to post-downgrade tasks.

## Updating SP firmware

If your storage system includes a Service Processor (SP), you must verify that it is running the correct firmware version and update it if it is not.

**Before you begin**

The reversion or downgrade process should be complete and the storage system should be running the target release.

## About this task

Data ONTAP software images include firmware for SP modules. If the firmware version on your SP module is outdated, you must update it before returning the reverted or downgraded system to production.

## Steps

1. Go to the system firmware information on the N series support site and determine the most recent firmware version for your SP module.
2. Enter the following command at the storage system CLI to determine the SP firmware version:

```
sp status
```

You see output similar to the following:

```
Service Processor Status: Online
Firmware Version: 1.2
...
```

If the SP firmware version in the command output is earlier than the most recent version on the N series support site, you must update your disk shelf firmware manually.

3. Click the `SP_FW.zip` link to download the file from the N series support site to your HTTP server.
4. At the storage system prompt, enter the following command:  

```
software update http://Web_server/SP_FW.zip -f
```
5. When the `software update` command is finished, enter the following command at the storage system prompt:  

```
sp update
```
6. When the system prompts you to update SP, enter **y** to continue.  
SP is updated and you are prompted to reboot SP. Wait approximately 60 seconds to allow SP to reboot.
7. Verify that the SP firmware has been updated by entering the following command:  

```
sp update
```
8. If the system is a partner in an HA pair, repeat Steps 4 through 7 on the partner system.

## Result

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.

## Completing post-downgrade tasks

After downgrading to an earlier Data ONTAP release, you might need to perform additional tasks.

You should verify that any services you halted manually restarted after the downgrade. If not, you should restart them manually and verify that any clients have appropriate access to storage system services.



# 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.

## 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 system upgrades and HA configuration upgrades when downtime is scheduled, but it 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.
- 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 15

[Service and protocol considerations](#) on page 130

[Updating firmware](#) on page 69

[Updating disk shelf firmware](#) on page 75

## 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.

## 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 compatibility and configuration information about FCP and iSCSI products, see the appropriate matrix at the N series Service and Support website at [www.ibm.com/storage/support/nseries/](http://www.ibm.com/storage/support/nseries/).

## 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.

## Understanding background disk firmware updates

When a storage system 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 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 is 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*.

---

## Copyright and trademark information

Copyright ©1994 - 2011 NetApp, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 2011 IBM Corporation. All rights reserved.

US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

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.

References in this documentation to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's or NetApp's intellectual property rights may be used instead of the IBM or NetApp product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM and NetApp, are the user's responsibility.

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

IBM, the IBM logo, and [ibm.com](http://www.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 <http://www.ibm.com/legal/copytrade.shtml>

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, FAServer, FastStak, FilerView, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service

Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

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.

---

## Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe on any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, N.Y. 10504-1785  
U.S.A.

For additional information, visit the web at:  
<http://www.ibm.com/ibm/licensing/contact/>

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

**INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.** Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM web sites are provided for convenience only and do not in any manner serve as an endorsement of those web sites. The materials at those web sites are not part of the materials for this IBM product and use of those web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.



# Index

- A**
- ACP firmware
    - updating 80
- B**
- BMC firmware 88
- C**
- CIFS
    - requires disruptive upgrade 20
  - CPU utilization, nondisruptive upgrade requirements 21
- D**
- Data ONTAP
    - downgrading an HA pair (disruptive) 125
    - downgrading an HA pair (nondisruptive) 123
    - downgrading deduplicated volumes 120
    - preparing for the upgrade 27
    - requirements for downgrading from the 8.0 release family 118, 119
    - requirements for reverting from the 8.0 release family 96, 97
    - upgrading a high-availability configuration (disruptive) 60
    - upgrading a high-availability configuration from an earlier release family nondisruptively 53
    - upgrading a high-availability configuration within a release family (nondisruptive) 58
    - upgrading a single system 63
  - Data ONTAP 8.0
    - reversion issues with VLANs 105
  - Data ONTAP software images
    - copy software images without installing 38
    - copying from a UNIX client 39
    - copying from a Windows client 40
    - getting from IBM 41
  - Data ONTAP system files
    - copying the software image to the HTTP server 38
    - downloading from IBM 39
    - installation overview 43
    - installation procedure for HTTP 43
      - installation procedure from /etc/software 47
      - managing from an HTTP server 37
      - managing with the software command 42
  - deduplication
    - upgrade requirements 28
  - disk firmware updates
    - about 73
  - disk firmware upgrades
    - background 131
  - disk shelf
    - hot-adding 75
  - disk shelf firmware updates
    - about 75
    - determining firmware versions 78
    - NDU requirements 76
  - disk shelf firmware upgrades
    - manual update procedure 78
  - disk utilization, nondisruptive upgrade requirements 21
  - disruptive system firmware update 72
  - DNS, enable 29
  - domain account, verifying 29
  - downgrade issues
    - compression for SnapMirror transfers 119
  - downgrading
    - supported scenarios 117
  - downgrading from Data ONTAP 8.0
    - VLAN configuration issues 105
  - downgrading to a previous release
    - downgrading from the 8.0 release family 118, 119
- F**
- firmware
    - ACP 80
  - firmware updates
    - disk shelf 75
  - firmware upgrades
    - BMC 88
    - disk 73
    - Flash Cache 93
    - PAM 93
    - RLM 82
    - SP (Service Processor) 81
    - system 69
  - Flash Cache firmware 93
  - FlexVol volumes

- nondisruptive upgrade requirements 21

## L

- LUN restore 99

## M

- major
  - nondisruptive upgrades 21
- minor
  - nondisruptive upgrades 21
- module firmware
  - disk shelf 75

## N

- NDU requirements 76
- nondisruptive upgrades
  - about 20
  - Data ONTAP software 29
  - preparing 29
  - requirements 21
  - system firmware 70
  - when not to use 20

## O

- out-of-order frame delivery, reverting with 114

## P

- PAM firmware 93
- post-upgrade tasks 65

## R

- release families
  - differentiating among 18
  - overview 18
  - upgrading between 18
  - upgrading within 19
- reversion issues
  - Compatibility issues attached with Brocade Fabric Operating System and Data ONTAP 107
  - out-of-order frame delivery 114
- reversion issues FlexClone files and LUNs, reverting
  - FlexClone files and LUNs 107
- revert
  - SnapMirror, preserve relationship 104

- reverting
  - supported scenarios 95
- reverting from Data ONTAP 8.0
  - enabling TOE 114
  - VLAN and base interface configuration issues 105
- reverting to a previous release
  - reverting from the 8.0 release family 96, 97
- RLM
  - firmware update problems, troubleshooting 87
  - troubleshooting firmware update problems 87
- RLM firmware 82
- rolling upgrade 20

## S

- SnapMirror
  - identifying destination volumes 51
  - issues for systems with synchronous SnapMirror 17
  - planning upgrades 17
  - revert, initialize 104
  - revert, preserve relationship 104
  - upgrade requirements 16
  - upgrade, preserve relationship 104
  - upgrading for volume replication 51
  - upgrading systems that are mirroring volumes to each other 18
- Snapshot copies
  - nondisruptive upgrade requirements 21
- software update command 43
- solid-state disks (SSDs)
  - reverting and 107
- SP (Service Processor) firmware 81
- space guarantees
  - reverting 103
  - reverting with 103
  - space guarantees enabled 103
- special system files
  - .bplusvtoc\_internal 99
  - .vtoc\_internal 99
- SSDs
  - reverting and 107
- storage download shelf command 78
- system firmware
  - about 69
  - disruptive firmware update procedure 72
  - nondisruptive upgrade 70
  - obtaining 69

**U**

## UNIX host

- mounting the system 39

## upgrade

- enabling DNS with Windows 2000 name addresses

  - 29

- maintaining service 129

- overview 13

- overview of requirements 13

- preparing for 27

- required intermediate upgrades 19

- resolving issues 23

- SnapMirror, preserve relationship 104

- system requirements 28

- verifying system domain account 29

- with session oriented protocols 131

- with stateless protocols 130

## upgrade host

- requirements 15

**V**

## VLAN configuration

- reversion issues with Data ONTAP 8.0 105

**W**

## Windows host

- mapping the root directory to a client share 41





NA 210-05294\_A0, Printed in USA

GA32-0724-05

