IBM System Storage N series

Clustered Data ONTAP 8.2 Upgrade and Revert/Downgrade Guide

SC27-6284-01

Contents

Preface

About this guide

This document applies to IBM N series systems running Data ONTAP, including systems with gateway functionality. If the terms *Cluster-Mode* or *clustered Data ONTAP* are used in this document, they refer to the Data ONTAP features and functionality designed for clusters, which are different from 7-Mode and prior Data ONTAP 7.1, 7.2, and 7.3 release families.

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 (accessed and navigated as described in *Websites* on page 7).

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 N series support website (accessed and navigated as described in *Websites* on page 7) 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 N series support website (accessed and navigated as described in *Websites* on page 7).

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 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, ensure that you run the latest level of firmware. Any firmware updates are posted to the N series support website (accessed and navigated as described in *Websites* on page 7).

> **Note:** If you do not see new firmware updates on the 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 email 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

Upgrading Data ONTAP clusters

Upgrading a cluster to the current Data ONTAP release requires preparation, the upgrade itself, and several post-upgrade procedures.

The upgrade process includes the following phases:

- Preparing for the upgrade
- Performing the upgrade
- Completing post-upgrade tasks

Plan to do the following:

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

  See the N series Interoperability Matrices website (accessed and navigated as described in *Websites* on page 7).

Unless otherwise indicated, the requirements and procedures in this guide apply to all supported:

- Data ONTAP 8.2.x platforms

  For more information about supported platforms, see the *Release Notes* for the target Data ONTAP release.
- Upgrade paths to and within the Data ONTAP 8.2 release family

  The supported upgrade paths include upgrades to releases in the Data ONTAP 8.2 release family from any 8.1.x release (*major upgrades*) and upgrades from 8.2.x to 8.2.z (*minor upgrades*).

Planning your upgrade

Because new features are introduced in each release of Data ONTAP, you must understand these features and their associated 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 cluster after the upgrade.
- Prepare 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 cluster before the upgrade.

Unless otherwise indicated, the requirements and procedures in this guide apply to all supported Data ONTAP 8.2.x platforms. For more information about supported platforms, see the *Release Notes* for this Data ONTAP release.

Related concepts

Reverting clusters to an earlier Data ONTAP release family on page 90
Downgrading clusters to an earlier release in the same release family on page 130

Types of cluster upgrades

Based on your requirements, you can upgrade a cluster to a new Data ONTAP release by performing a nondisruptive upgrade or a disruptive upgrade.

In a *nondisruptive upgrade* (NDU), the cluster remains online and continues to serve data during the upgrade. There are two methods for performing a NDU:

- Rolling upgrade

  In this method, a node is taken offline and upgraded while its partner takes over its storage. When the node upgrade is complete, the assisting partner gives control back to the original partner and the process is repeated, this time on the partner. Each additional HA pair is upgraded in sequence until all pairs are running the target release.

  You can use this method for a cluster of any size, but it is required for clusters consisting of 2-6 nodes.
- Batch upgrade

  In this method, the cluster is separated into two batches, each of which contains multiple HA pairs. In the first batch, one node in each HA pair is taken offline and upgraded while their partners take over their storage. When the upgrade is completed for the first half of all of the HA pairs, their partners give control back to them, and the process is repeated, this time on the partners. The process is then repeated on the second batch.

  You can only use this method if the cluster consists of 8 or more nodes, and if it is configured for network-attached storage (NAS) only.

In a *disruptive upgrade*, the cluster is taken offline to perform the upgrade. This type of upgrade involves disabling storage failover for each HA pair and then rebooting each node.

You can perform a disruptive upgrade on a cluster of any size.

Cluster upgrade checklist

You can use this checklist to record your progress as you prepare for the upgrade, perform the upgrade, and complete post-upgrade tasks.

Steps for preparing to upgrade

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

Condition	Complete?
Software and hardware support in the target release is confirmed. To confirm hardware support, visit *IBM System Storage N series Introduction and Planning Guide* at the N series support website (accessed and navigated as described in *Websites* on page 7)	
Cluster network and management network switches are supported. The software, firmware, and reference configuration files used by cluster and management Ethernet switches must be compatible with Data ONTAP. When planning a Data ONTAP deployment or upgrade, you must consult the *Cluster and Management Ethernet Switch Matrix* to determine whether updates are also required to the switch configurations.	
The SAN configuration is fully supported. All SAN components—including target Data ONTAP software version, host OS and patches, required Host Utilities software, and adapter drivers and firmware—should be listed with the compatibility and configuration information about FCP and iSCSI products. See the appropriate matrix at the N series Interoperability Matrices website (accessed and navigated as described in *Websites* on page 7).	
All release-specific upgrade issues have been resolved.	
You have clustershell access privileges.	
You have created a performance baseline. You use Perfstat Converged to create a performance baseline.	
A remote management device is configured for each node. You should have either a Service Processor (SP) or Remote LAN Module (RLM) device configured. For more information, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.	
The HA pair is verified. You can use the Config Advisor tool to check for common configuration errors.	
Deduplicated volumes and aggregates contain sufficient free space. Each deduplicated volume must contain at least 4% free space. Each aggregate that contains a deduplicated volume must contain at least 3% free space.	

Condition	Complete?
The cluster is in quorum. Ensure that all nodes are participating in a quorum and all rings are in the quorum. Verify also that the per-ring quorum master is the same for all nodes.	
The cluster and Storage Virtual Machines (SVMs) are running and healthy. All aggregates and volumes should be healthy and online before proceeding with the upgrade. You can use the `cluster show` command to verify the status of the nodes.	
LIFs are online and on their correct home ports. You can use the `network interface` command to display and modify LIF configuration.	
The LIF failover configuration is verified. Each data LIF must have the correct failover policy, failover group, and failover targets applied.	
Automatic LIF rebalancing is disabled (for batch upgrades only).	
Any CIFS sessions that are not continuously available have been terminated.	
The system time is synchronized across the cluster.	
Each node is running Data ONTAP 8.1 or later. Ensure that you are running on the minimum software version allowed for the upgrade by running the `system node image show` command.	
You have the target Data ONTAP software image available on an HTTP server. Download the software image for the target Data ONTAP release from the N series support website (accessed and navigated as described in *Websites* on page 7), and place it on an HTTP server that is accessible by each node.	
The target Data ONTAP software images are installed on each node and set as the alternate boot device image. You can use the `system node image update` command to install the software images. You can use the `system node image show` command to verify that the software images are installed as the alternate boot image on each node.	

Condition	**Complete?**
SnapMirror operations are suspended.	
No jobs are running. If any aggregate, volume, mirror, NDMP (dump or restore), or Snapshot jobs (such as create, delete, move, modify, replicate, and mount jobs) are running or queued, allow the jobs to complete successfully or stop the queued entries.	

Steps for performing a rolling upgrade

If you are performing a nondisruptive upgrade by using the rolling upgrade method, the upgrade is complete when all of the following steps have been completed:

Condition		**Complete?**
You have verified that the cluster is ready to be upgraded nondisruptively.		
The first HA pair is upgraded.	The first node in the HA pair is upgraded.	
	The node's partner is upgraded.	
	You have verified that the HA pair was upgraded successfully.	
If needed, the second HA pair is upgraded.	The first node in the HA pair is upgraded.	
	The node's partner is upgraded.	
	You have verified that the HA pair was upgraded successfully.	
If needed, the third HA pair is upgraded.	The first node in the HA pair is upgraded.	
	The node's partner is upgraded.	
	You have verified that the HA pair was upgraded successfully.	
You have verified that the cluster upgrade was completed successfully.		

Steps for performing a batch upgrade

If you are performing a nondisruptive upgrade by using the batch upgrade method, the upgrade is complete when all of the following steps have been completed:

Condition		**Complete?**
You have verified that the cluster is ready to be upgraded nondisruptively.		
The cluster is divided into two batches. If the cluster has an even number of HA pairs, then each batch should contain half of the HA pairs. If the cluster has an odd number of HA pairs, then the first batch should contain one more HA pair than the second batch.		
The first batch is upgraded.	The first node in each HA pair is upgraded.	
	The nodes' partners are upgraded.	
	You have verified that the HA pairs were upgraded successfully.	
The second batch is upgraded.	The first node in each HA pair is upgraded.	
	The nodes' partners are upgraded.	
	You have verified that the HA pairs were upgraded successfully.	
You have verified that the cluster upgrade was completed successfully.		

Steps for performing a disruptive upgrade

If you are performing a disruptive upgrade, the upgrade is complete when all of the following steps have been completed:

Condition	**Complete?**
You have verified that the cluster is ready to be upgraded disruptively.	
Storage failover is disabled and each node is rebooted.	
You have verified that the cluster upgrade was completed successfully.	

Steps for after upgrading

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

Condition	**Complete?**
The cluster and SVMs are running and healthy. All aggregates and volumes should be healthy and online before proceeding with the upgrade. You can use the `cluster show` command to verify the status of the nodes.	

Condition	Complete?
LIFs are online and on their correct home ports. You can use the `network interface` command to display and modify LIF configuration.	
A namespace mirror constituent exists for an Infinite Volume. If the cluster contains an Infinite Volume that spans two or more nodes, you must create a namespace mirror constituent to provide data protection for the namespace constituent.	
The cluster management LIF is set correctly for Remote Support Agent. You can set the cluster management LIF by using the `rsa setup` command on each SP or RLM device.	
SnapMirror operations are resumed. If you suspended SnapMirror operations prior to performing a nondisruptive upgrade, you should resume the operations after the upgrade is complete.	
Older style data protection mirror relationships are updated. SnapMirror relationships that you have before upgrading to Data ONTAP 8.2 must use the Data ONTAP 8.2 syntax style if you want to make use of its features and improvements.	
Automatic LIF rebalancing is reenabled.	

Related concepts

Identifying potential upgrade issues on page 24

Performing the software upgrade on page 46

Related tasks

Preparing for the Data ONTAP cluster upgrade on page 27

Completing post-upgrade tasks for cluster upgrades on page 73

Related information

IBM N series Interoperability Matrix: www.ibm.com/systems/storage/network/interophome.html

Cluster upgrade requirements

There are release and configuration requirements that your cluster should meet before you perform an upgrade. Additionally, there are mixed version requirements that you should be aware of while you are performing the upgrade.

Release requirements

You can upgrade clusters to the Data ONTAP 8.2 release family from any 8.1.x release. Minor upgrades within the Data ONTAP 8.2 release family are supported from all previous 8.2 releases.

Note: If you are running Data ONTAP GX 10.x, do not attempt to upgrade clusters to the Data ONTAP 8.2 release family on your own; doing so is an unsupported operation. Contact your IBM representative for assistance.

If you want to upgrade to a Data ONTAP 8.2 release from a release earlier than 8.1, you must perform an intermediate upgrade (also known as a multi-hop upgrade) to the latest Data ONTAP 8.1 release before upgrading to the target Data ONTAP 8.2 release.

Configuration requirements

The cluster must meet the following configuration requirements before you upgrade it:

- Because failed disk drives prevent giveback operations and can introduce loop instability throughout the cluster, you must remove or replace all failed disk drives *before* beginning the upgrade process.

  For more information about identifying and removing failed disks, see the *Clustered Data ONTAP Physical Storage Management Guide*.
- If your cluster serves NFS clients, you must use hard mounts.

  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.
- If your cluster serves CIFS clients, nondisruptive upgrades are supported for Hyper-V over SMB solutions.

  Hyper-V over SMB solutions enable Hyper-V and the contained virtual machines to remain online and to provide continuous availability during the Data ONTAP upgrade. For more information and configuration limits and requirements, see the *Clustered Data ONTAP File Access Management Guide for CIFS*.

  For all other CIFS configurations, client sessions are terminated. You should direct users to end their sessions before you upgrade to prevent data loss.
- Upgrades might be disruptive if the cluster is actively serving Network Data Management Protocol (NDMP) clients that cannot be postponed.

  Because this protocol is session-oriented, current sessions must finish, and the service must be disabled to use nondisruptive upgrades.

Mixed version requirements

Data ONTAP clusters can operate for a limited time in a *mixed version* state, in which nodes in a cluster are running Data ONTAP versions from different release families. However, the upgrade is not complete until all nodes are running the new target release.

When the cluster is in a mixed-version state, you should not enter any commands that alter the cluster operation or configuration except as necessary to satisfy upgrade requirements; monitoring operations are permitted.

Related tasks

Creating a performance baseline with Perfstat Converged on page 29

Cluster upgrade limits

Before performing an upgrade, you must verify that your cluster does not exceed the platform system limits. SAN and Hyper-V over SMB also have limits that you should verify in addition to the platform system limits.

General platform limits

You should not exceed the system limits for your platform. To obtain your platform's system limits, see the *IBM System Storage N series Introduction and Planning Guide* at the N series support website (accessed and navigated as described in *Websites* on page 7). In addition, you should not exceed the maximum values for the following system elements:

Element	Value (per node)	Command to display values
Snapshot copies	20,000	`volume snapshot show`
CPU utilization *	No greater than 50%	`node run -node` *`node_name`* `-command sysstat -c 10 -x 3`
Disk utilization *	No greater than 50%	

* Before upgrading Data ONTAP, you should monitor CPU and disk utilization for 30 seconds. The values in the `CPU` and `Disk Util` columns should not exceed 50% for all ten measurements reported. You should ensure that no additional load is added to the cluster until the upgrade completes.

SAN limits

If your cluster is configured for SAN, you should not exceed the limits defined in the *Clustered Data ONTAP SAN Configuration Guide*.

Hyper-V over SMB limits

If you have configured a Hyper-V over SMB solution, you should not exceed the maximum values for the following system elements on all platforms:

Element	Value (per node)	Command to display values
FlexVol volumes	200	`volume show -node` *`node_name`*
LUNs	400	`lun show`
Snapshot copies	2,000	`volume snapshot show`

Required upgrade sequence for a batch upgrade

Performing a batch upgrade requires you to upgrade sets of nodes in a particular sequence. By following the required sequence, you can ensure that the cluster will remain up and continue serving data while you perform the NDU.

The following diagram illustrates the batch upgrade sequence for a 12-node cluster:

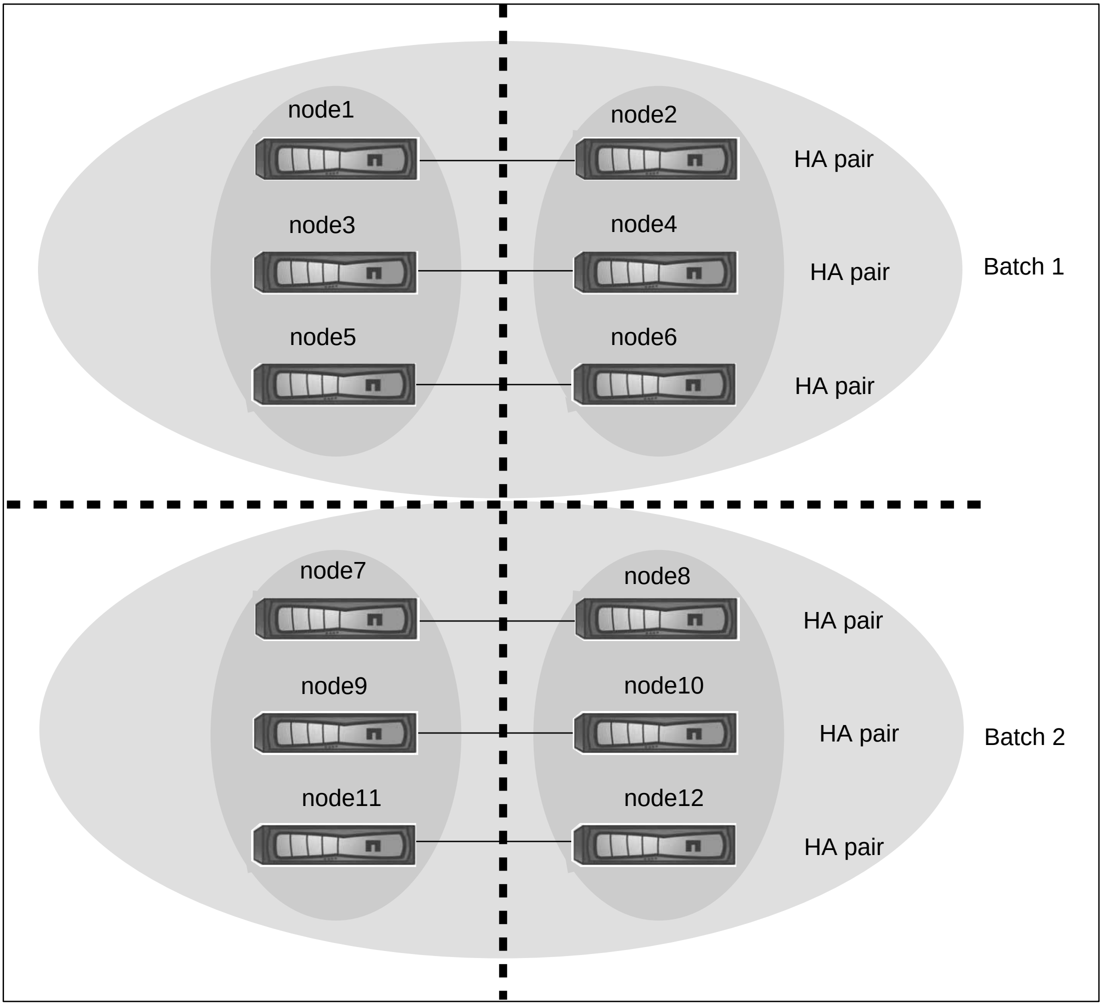

Note: The node names in this diagram are provided for example purposes only. The nodes are numbered to depict the proper upgrade sequence.

This cluster is divided into two batches, each containing an equal number of HA pairs. The nodes are then upgraded in the following sequence:

1. Batch 1:
    a. Upgrade node1, node3, and node5 concurrently
    b. Upgrade node2, node4, and node6 concurrently
2. Batch 2:
    a. Upgrade node7, node9, and node11 concurrently
    b. Upgrade node8, node10, and node12 concurrently

If a cluster contains an odd number of HA pairs, the first batch should contain the larger number of HA pairs. For example, the following diagram shows the batch upgrade sequence for a 14-node cluster, in which the first batch contains 4 HA pairs, and the second batch contains 3 HA pairs:

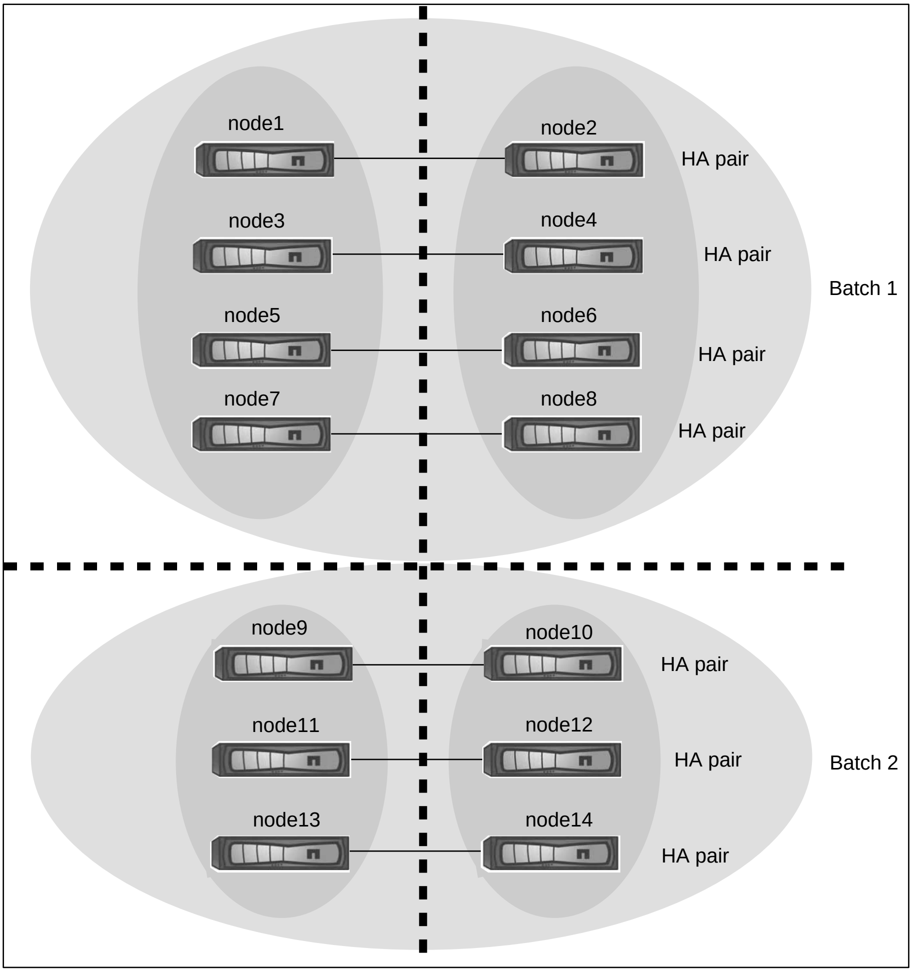

In this 14-node cluster, the nodes are upgraded in the following sequence:

1. Batch 1:
    a. Upgrade node1, node3, node5, and node7 concurrently
    b. Upgrade node2, node4, node6, and node8 concurrently
2. Batch 2:
    a. Upgrade node9, node11, and node13 concurrently
    b. Upgrade node10, node12, and node14 concurrently

Guidelines for estimating the duration of the upgrade process

For each HA pair, you should plan for approximately 30 minutes to complete preparatory steps, 60 minutes to perform the upgrade, and 30 minutes to complete post-upgrade steps.

The batch method for performing a nondisruptive upgrade enables you to upgrade all of the HA pairs in a batch concurrently. Accordingly, if you perform a batch upgrade, the total amount of time required to upgrade the cluster should be similar to the amount of time required to upgrade two HA pairs regardless of the size of the cluster.

The upgrade duration guidelines are based on typical configurations and workloads. You can use these guidelines to estimate the time it will take to perform a nondisruptive upgrade in your environment.

Upgrade considerations for an Infinite Volume

When upgrading a cluster that is running Data ONTAP 8.1.1 or later, and the cluster contains an Infinite Volume, you must understand the order of upgrading the nodes and the procedures you must perform after the upgrade.

If technical support created a data protection mirror copy for the namespace constituent in an Infinite Volume running Data ONTAP 8.1.x, you should upgrade the node with the data protection mirror copy for the namespace constituent first. If technical support did not create a data protection mirror copy for the namespace constituent, you can upgrade the nodes in any order.

After the upgrade, Infinite Volumes that span two or more nodes must have a namespace mirror constituent to provide data protection for the Infinite Volume's namespace constituent. The upgrade process does not automatically create the namespace mirror constituent. You must review the post-upgrade procedure for Infinite Volumes and plan for how to create a namespace mirror constituent after the upgrade.

After the upgrade, Infinite Volumes that are in a data protection mirror relationship must use Storage Virtual Machine (SVM) peering between the SVM with Infinite Volume in the source cluster and the SVM with Infinite Volume in the destination cluster to enable the data protection features available in Data ONTAP 8.2.x. The upgrade process does not automatically configure SVM peering or the new relationship capabilities. You must review the post-upgrade procedure and plan for how to upgrade the data protection mirror relationship to use SVM peering and the new relationship capabilities after the upgrade.

Related tasks

Creating a namespace mirror constituent for an upgraded Infinite Volume on page 79
Upgrading older style data protection mirror relationships on page 82

Upgrade requirements for SnapMirror

If you are upgrading a cluster that is running SnapMirror, you must suspend SnapMirror operations for each node that contains destination volumes.

For SnapMirror volume replication, the destination node must use a Data ONTAP version that is equal to or later than that of the SnapMirror source node. To prevent SnapMirror transfers from failing, you must suspend SnapMirror operations and, in some cases, upgrade destination nodes before upgrading source nodes.

Suspending SnapMirror operations for the duration of the NDU

The simplest method for upgrading in a SnapMirror environment is to suspend all SnapMirror operations, perform the upgrade, and then resume the SnapMirror operations. However, no SnapMirror transfers will occur during the entire NDU. You must use this method if you are performing a batch upgrade or if your cluster contains nodes that are mirroring volumes to each other.

Suspending SnapMirror operations one destination volume at a time

Alternatively, you can suspend SnapMirror transfers for a particular destination volume, upgrade the node (or HA pair) that contains the destination volume, upgrade the node (or HA pair) that contains the source volume, and then resume the SnapMirror transfers for the destination volume. By using this method, SnapMirror transfers for all other destination volumes can continue while the nodes that contain the original destination and source volumes are upgraded.

This method requires you to perform a rolling upgrade; batch upgrades are not permitted.

For more information about running SnapMirror on clusters configured for network-attached storage (NAS), see the *Clustered Data ONTAP Data Protection Guide*.

Information about low-memory gateway upgrade

While upgrading your low-memory gateway (a storage system with physical memory of 8 GB or less; for example V3140, V3160, V3210, and V3240) from a previous Data ONTAP release to Data ONTAP 8.2.1, your system might undergo multiple reboots.

Multiple reboots occur if the number of storage devices discovered by the gateway exceeds the maximum number of storage devices allowed for that platform. This happens because, in Data ONTAP 8.2.1, memory is allocated dynamically for storage devices as you present the storage devices to your gateway.

For information about the number of storage devices allowed for different platforms, see the *IBM System Storage N series Introduction and Planning Guide* at the N series support website (accessed and navigated as described in *Websites* on page 7).

Identifying potential upgrade issues

Every Data ONTAP release family has unique upgrade requirements that you must understand before you decide to upgrade.

Before you upgrade, you must 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.

Be sure to consult the *Release Notes* for the upgrade target release for a complete listing of upgrade issues.

Upgrade issues with the Data ONTAP 8.2 release family

You must understand and resolve any known technical issues before you upgrade a cluster to Data ONTAP 8.2 and later releases.

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

- Starting with Data ONTAP 8.2.1, there is a change in how Data ONTAP handles file names containing UTF-16 supplementary characters that you must be aware of if your environment uses such file names.
- After you upgrade to Data ONTAP 8.2, the license-list-info ZAPI is no longer supported.

  As a result, you must install Data ONTAP management software versions that are supported by Data ONTAP 8.2.
- During a Data ONTAP upgrade, LUNs are assigned new revision numbers.

  Windows Server 2008 and Windows Server 2012 interpret the LUNs with new revision numbers as new disks and set them offline after a host reboot; this status is shown in Windows management interfaces after the upgrade. Windows Server 2003 ignores LUN revision numbers.
- Stale deduplication-related metadata can exist in the FlexVol volumes and aggregates on your system, resulting in slow deduplication processing or your systems running out of space.

  Your systems can experience this problem if you are upgrading from a Data ONTAP 8.1 release prior to 8.1.2P4, if deduplication is enabled on any FlexVol volume, and if a FlexVol volume or the associated aggregate is more than 70 percent full.
- Starting with Data ONTAP 8.2, all license keys are 28 characters in length.

  Licenses installed prior to Data ONTAP 8.2 continue to work after you upgrade to Data ONTAP 8.2 or later. However, if you need to reinstall a license when you are running Data ONTAP 8.2 or later, the old key is not accepted.
- Starting with Data ONTAP 8.2, a baseline SP firmware image is packaged with the Data ONTAP image.

By default, the SP automatic update functionality is enabled. You have the option to manually trigger an SP update.

- Starting with Data ONTAP 8.1.2, the software install or upgrade includes new disk shelf firmware versions for EXN1000 unit and EXN3000 disk shelves that provide enhanced disk error detection and prediction capabilities.

  Therefore, after upgrading to this version of Data ONTAP, you might experience an increase in the number of disk failures for certain disk shelf and disk models.
- If you have a N6210 or N6040 system with Flash Cache modules installed, do not upgrade to clustered Data ONTAP 8.2 or later 8.2.x releases.

  Flash Cache modules are not supported on N6210 or N6040 systems running clustered Data ONTAP 8.2 or later 8.2.x releases.
- Starting with Data ONTAP 8.2.1, the off-board antivirus feature is supported.

  To enable the off-board antivirus feature, you must ensure that all the nodes in the cluster are running Data ONTAP 8.2.1 or later.
- Beginning with Data ONTAP 8.2, striped volumes are no longer supported.

  You should not upgrade to Data ONTAP 8.2 if there are striped volumes in your environment.
- Beginning with Data ONTAP 8.2, NFSv2 is no longer supported and all NFSv2 functionality has been removed from Data ONTAP.

  If you still require the use of the NFSv2 protocol in your environment, do not upgrade to Data ONTAP 8.2 or later until you have discontinued use of NFSv2.
- If you have a flex_clone license, you must turn on the `licensed_feature.flex_clone` option after you upgrade to Data ONTAP 8.2.
- In Data ONTAP 8.2.x releases, you must ensure that there is sufficient space on the aggregate and on the volume to upgrade the deduplication metadata.

  You can use the `df` command to check the space available on the aggregate and on the volume.
- Modifying certain 7-Mode CIFS-specific options to a non-default value on a controller running clustered Data ONTAP 8.1.x or 8.2.x and then upgrading to a later version of 8.2.x might lead to recursive service disruptions.
- Starting with Data ONTAP 8.1.3, support for on-board antivirus is discontinued.

  Therefore, you must disable on-board antivirus and unschedule all the antivirus on-demand (AVOD) jobs before upgrading to Data ONTAP 8.2 and later.
- Beginning with Data ONTAP 8.2.1, Data ONTAP now sends DNS queries to resolve domain names in export policy rules by default over the data LIF, not the node management LIF.
- Beginning with Data ONTAP 8.2.1, Data ONTAP now sends DNS queries to resolve host names in export policy rules by default over the data LIF, not the node management LIF.
- Exceeding the supported limits for the number of CIFS local users and groups on the cluster might lead to a disruption of service.

  You must not exceed the supported limits on the number of CIFS local users and groups that you can create on the cluster.
- When NFSv3 or NFSv4.x clients attempt to modify file attributes using a `chmod`, `chown`, or `chgrp` command for a path that represents a junction point on volumes with UNIX security style, the operation appears successful but actually fails silently.

  If the client attribute cache is enabled, it is populated with incorrect information.

- If you upgrade from a version of Data ONTAP that does not support SMB 3.0 to a version of Data ONTAP that does, reconnections to mapped SMB shares from Windows 8 or Windows 2012 Server might fail with an “Invalid Signature” error.

  Data ONTAP 8.2 and later releases support SMB 3.0.

Changes to behavior in the Data ONTAP 8.2 release family

You should be aware of changes in Data ONTAP behavior that occur if you upgrade a cluster to Data ONTAP 8.2 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.

- Starting with Data ONTAP 8.2, the Network Time Protocol (NTP) is always enabled on the cluster.

  If you manually set the time, the setting takes effect on all nodes in the cluster.
- Starting with Data ONTAP 8.2, switchless two-node clusters of N3150, N3220 and N3240 systems can be configured using only a single 10-GbE port on each node for the cluster network.
- Starting with Data ONTAP 8.2, Element Manager (including ClusterView) is no longer available.

  You can use the Data ONTAP command-line interface or OnCommand System Manager, a web-based graphical management interface, to manage the cluster.
- Starting in Data ONTAP 8.2, aggregate Snapshot copy automatic deletion is always enabled and cannot be disabled.
- In previous versions of Data ONTAP, you could set the fractional reserve setting for a FlexVol volume to any value between 0 and 100, inclusive.

  Starting with Data ONTAP 8.2, the fractional reserve setting can be set only to 0 or 100.
- Starting with Data ONTAP 8.2, Data ONTAP generates aliases based on SN instead of the WWN.
- In the node-scoped NDMP mode, you must use NDMP-specific credentials to access a storage system in order to perform tape backup and restore operations.
- Before using NFSv4.1 or pNFS on your storage system with NFS clients, you should take certain precautions to ensure successful deployment in your environment.
- When upgrading to Data ONTAP 8.2, the cluster can operate in a mixed-version state, in which some nodes are running release 8.2 and some are running release 8.1.

  While in this state, the cluster switch health monitor might generate health alerts. These alerts trigger AutoSupport messages that include an incorrect subject.
- The default LDAP client schema AD-SFU that was provided by Data ONTAP 8.1 does not work with Windows Services for UNIX (SFU) or Identity Management for UNIX (IDMU).

  During the upgrade to Data ONTAP 8.2, this problem is corrected.
- SnapMirror relationships that you had before upgrading to Data ONTAP 8.2 must use the Data ONTAP 8.2 syntax style if you want to make use of its features and improvements.
- For an intercluster Storage Virtual Machine (SVM) peer relationship that is in pending state, if the cluster administrator of the local cluster deletes the SVM peer relationship and the cluster

administrator of the peer cluster accepts the SVM peer relationship simultaneously, then both the clusters might have an inconsistent state with respect to SVM peer relationship.

- Avoid configuring LIFs with addresses in the 192.168.1/24 and 192.168.2/24 subnets.

  LIFs configured with these addresses may conflict with private iWARP interfaces and might result in the LIFs failing to come on line after a node reboot or a LIF migration.
- Fast path is enabled by default in all nodes in the cluster.

  However, you might have to disable fast path in certain scenarios to avoid issues such as performance degradation or failed software upgrade.
- If you are configuring volumes with zero fractional reserve and you are using certain technologies or Data ONTAP features, you must take some extra precautions to avoid out of space errors.
- If you configure ACP to use a different port than e0P (the default) on N3220 and N3240 systems, the internal ACPP module IOM6E becomes unresponsive, disabling ACP for the local SAS expander.

  Unless you are experiencing hardware issues on port e0P, always use the default port for ACP for the N3220 and N3240 systems.
- FlexCache volumes are not supported on the following storage systems running clustered Data ONTAP 8.2: N6040, N6060, N6210, and N6240.
- Starting with Data ONTAP 8.2, the new NFS server parameter `-mount-rootonly` is enabled by default, meaning Data ONTAP by default now rejects all NFS mount requests from nonreserved ports (ports numbered 1024 or higher).

  If your environment has clients that require mounting NFS exports using nonreserved ports, you must manually disable this parameter for each SVM after upgrading to Data ONTAP 8.2 or later.

Preparing for the Data ONTAP cluster upgrade

Before installing the latest Data ONTAP release on your cluster, you must perform several cluster health checks, including ensuring that the cluster is running and healthy, verifying that the cluster is in a quorum, and verifying the Storage Virtual Machine (SVM) health.

Verifying the HA pair cabling and configuration

You can use the Config Advisor tool to check for common configuration errors.

About this task

Config Advisor is a configuration validation and health check tool for IBM N series systems. It can be deployed at both secure sites and non-secure sites for data collection and system analysis.

Note: Support for Config Advisor is limited, and available only online.

Steps

1. To obtain the Config Advisor tool, contact your technical support representative.

2. After running Config Advisor, review the tool's output and follow the recommendations to address any issues discovered.

Verifying that deduplicated volumes and aggregates contain sufficient free space

Before upgrading Data ONTAP, you must verify that any deduplicated volumes, and the aggregates that contain them, have sufficient free space for the deduplication metadata. If there is insufficient free space, deduplication will be disabled when the Data ONTAP upgrade is completed.

About this task

Each deduplicated volume must contain at least 4% free space. Each aggregate that contains a deduplicated volume must contain at least 3% free space.

Steps

1. Determine which volumes are deduplicated:

`volume show -is-sis-volume true`

Example

This example displays a deduplicated volume and the aggregate that contains it.

```
cluster1::> volume show -is-sis-volume true
Vserver   Volume       Aggregate    State      Type       Size  Available Used%
--------- ------------ ------------ ---------- ---- ---------- ---------- -----
vs1       vol_2        aggr_2       online     RW         20GB    18.74GB    6%
```

2. Determine the free space available on each volume that you identified:

`df -vserver` ***`Vserver_name`*** **`-volume`** ***`volume_name`***

Each deduplicated volume must not contain more than 96% used capacity.

Example

In this example, the `capacity` field displays the free space available on the deduplicated volume identified earlier (vol_2).

```
cluster1::> df -vserver vs2 -volume vol_2
Filesystem              kbytes       used      avail capacity  Mounted on
/vol/vol_2/           19456000     264000   19192000       1%  /
/vol/vol_2/.snapshot      1024          0       1024       0%  //.snapshot
2 entries were displayed.
```

For details on how to increase the size of a volume, see the *Clustered Data ONTAP Logical Storage Management Guide*.

3. Identify the free space available on each aggregate that contains a deduplicated volume:

df -A -aggregate *aggregate_name*

Each aggregate must not contain more than 97% used capacity.

Example

In this example, the `capacity` field displays the free space available on the aggregate containing the deduplicated volume (aggr_2).

```
cluster1::> df -A -aggregate aggr_2
Aggregate                kbytes       used      avail capacity
aggr_2                344220000   20944000  323276000       6%
aggr_2/.snapshot              0          0          0       0%
2 entries were displayed.
```

For details on how to increase the size of an aggregate, see the *Clustered Data ONTAP Physical Storage Management Guide*.

Creating a performance baseline with Perfstat Converged

The Performance and Statistics Collector (Perfstat Converged) is a cluster diagnostics data collection tool, available on the N series support website (accessed and navigated as described in *Websites* on page 7), that enables you to establish a performance baseline for comparison after the upgrade. You should create a Perfstat report before upgrading.

Before you begin

The diag user account must be unlocked. For details about unlocking this account, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

About this task

You should create a Perfstat report during a typical usage time; this takes about 30 minutes.

Steps

1. Download Perfstat Converged from the N series support website (accessed and navigated as described in *Websites* on page 7).
2. Enter the following command during a typical usage period:

   perfstat8 *cluster_management_IP_address* -m c -t 4 -i 5 -z

After you finish

You should retain the output file for several weeks after the Data ONTAP upgrade is complete.

Verifying that the cluster is in quorum

Before and after you perform an upgrade, reversion, or downgrade, you must ensure that all nodes are participating in a replicated database (RDB) quorum and that all rings are in the quorum. You must also verify that the per-ring quorum master is the same for all nodes.

About this task

For more information about cluster replication rings and RDB quorums, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

Steps

1. Set the privilege level to advanced:
   `set -privilege advanced`

   Enter **`y`** to continue.
2. Display each RDB process:

To display this RDB process...	Enter this command...
Management application	**`cluster ring show -unitname mgmt`**
Volume location database	**`cluster ring show -unitname vldb`**
Virtual-Interface manager	**`cluster ring show -unitname vifmgr`**
SAN management daemon	**`cluster ring show -unitname bcomd`**

Example

This example shows the volume location database process for a cluster running Data ONTAP 8.1.x:

```
cluster1::*> cluster ring show -unitname vldb

Node   UnitName Epoch DB Epoch DB Trnxs Master
------ -------- ----- -------- -------- ---------
node0  vldb     154   154      14847    node0
node1  vldb     154   154      14847    node0
node2  vldb     154   154      14847    node0
node3  vldb     154   154      14847    node0
4 entries were displayed.
```

Example

This example shows the volume location database process for a cluster running Data ONTAP 8.2.x:

```
cluster1::*> cluster ring show -unitname vldb
Node      UnitName Epoch    DB Epoch DB Trnxs Master    Online
```

```
--------- -------- -------- -------- -------- --------- ---------
node0     vldb     154      154      14847    node0     master
node1     vldb     154      154      14847    node0     secondary
node2     vldb     154      154      14847    node0     secondary
node3     vldb     154      154      14847    node0     secondary
4 entries were displayed.
```

For each process, verify the following configuration details:

- The relational database epoch and database epochs match for each node.
- The per-ring quorum master is the same for all nodes.
  Note that each ring might have a different quorum master.

3. Return to the admin privilege level:

`set -privilege admin`

4. If you are operating in a SAN environment, verify that each node is in a SAN quorum:

`event log show -messagename scsiblade.*`

The most recent `scsiblade` event message for each node should indicate that the scsi-blade is in quorum. During the upgrade, reversion, or downgrade process, each node will temporarily fall out of SAN quorum. Therefore, if you are verifying the SAN quorum after completing an upgrade, reversion, or downgrade, you may notice critical event messages warning you that the nodes were previously out of SAN quorum.

If a node is out of SAN quorum, you can use the `storage failover takeover` and `storage failover giveback` commands to perform a planned takeover and giveback with the node's high-availability partner and bring the node back into SAN quorum.

Example

In this example, both nodes in the cluster are in SAN quorum.

```
cluster1::> event log show -messagename scsiblade.*
Time                Node             Severity      Event
------------------- ---------------- ------------- ---------------------------
8/13/2013 14:03:51  node0            INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
8/13/2013 14:03:51  node1            INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
```

Example

This example shows a two-node cluster after performing an upgrade. Each node shows a previous out of SAN quorum event message from when the node was upgraded. However, the most recent event message for each node shows that both nodes are presently in SAN quorum.

```
cluster1::> event log show -messagename scsiblade.*
Time                Node             Severity      Event
------------------- ---------------- ------------- ---------------------------
8/13/2013 15:37:51  node1            INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
8/13/2013 15:31:26  node1            CRITICAL      scsiblade.out.of.quorum: This node ...
8/13/2013 15:31:26  node1            INFORMATIONAL scsiblade.ics.trace.dumpfile: The ...
8/13/2013 15:31:26  node1            INFORMATIONAL scsiblade.ics.trace.dumpfile: The ...
8/13/2013 15:30:43  node0            INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
8/13/2013 15:24:16  node0            CRITICAL      scsiblade.out.of.quorum: This node ...
8/13/2013 15:24:16  node0            INFORMATIONAL scsiblade.ics.trace.dumpfile: The ...
8/13/2013 15:24:16  node0            INFORMATIONAL scsiblade.ics.trace.dumpfile: The ...
```

Verifying cluster and SVM health

Before and after you upgrade, revert, or downgrade a cluster, you should verify that the nodes are healthy and eligible to participate in the cluster, and the aggregates and volumes are online.

Steps

1. Verify that the nodes in the cluster are online and are eligible to participate in the cluster:

   `cluster show`

   #### Example

   ```
   cluster1::> cluster show
   Node                  Health  Eligibility
   --------------------- ------- ------------
   node0                 true    true
   node1                 true    true
   ```

   If any node is unhealthy or ineligible, check EMS logs for errors and take corrective action. For more information about EMS messages, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

2. Determine if any disk drives are broken, undergoing maintenance, or reconstructing:

To check for...	Do this...
Broken disks	**a.** Display any broken disks: **`storage disk show -state broken`** **b.** Remove or replace any broken disks.
Disks undergoing maintenance or reconstructing	**a.** Display any disks in maintenance, pending, or reconstructing states: **`storage disk show -state maintenance|pending|reconstructing`** **b.** Wait for the maintenance or reconstruction operation to complete before proceeding.

3. To verify that all aggregates are online, display the state of physical and logical storage, including storage aggregates:

   `storage aggregate show -state !online`

   This command displays the aggregates that are *not* online.

   #### Example

   ```
   cluster1::> storage aggregate show -state !online
   There are no entries matching your query.
   ```

For more information about managing aggregates, see the *Clustered Data ONTAP Physical Storage Management Guide*.

4. To verify that all volumes are online, display any volumes *not* online:

   `volume show -state !online`

   #### Example

   ```
   cluster1::> volume show -state !online
   There are no entries matching your query.
   ```

   For more information about managing volumes, see the *Clustered Data ONTAP Logical Storage Management Guide*.

Enabling and reverting LIFs to home ports

During a reboot, some LIFs might have been migrated to their assigned failover ports. Before and after you upgrade, revert, or downgrade a cluster, you must enable and revert any LIFs that are not on their home ports.

About this task

The `network interface revert` command reverts a LIF that is not currently on its home port back to its home port, provided that the home port is operational. A LIF's home port is specified when the LIF is created; you can determine the home port for a LIF by using the `network interface show` command.

Steps

1. Display the status of all LIFs:

   `network interface show`

   #### Example

   This example displays the status of all LIFs for a Storage Virtual Machine (SVM, formerly known as Vserver).

   ```
   cluster1::> network interface show -vserver vs0
               Logical    Status     Network            Current       Current Is
   Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
   ----------- ---------- ---------- ------------------ ------------- ------- ----
   vs0
               data001    down/down  192.0.2.120/24     node0         e0e     true
               data002    down/down  192.0.2.121/24     node0         e0f     true
               data003    down/down  192.0.2.122/24     node0         e2a     true
               data004    down/down  192.0.2.123/24     node0         e2b     true
               data005    down/down  192.0.2.124/24     node0         e0e     false
               data006    down/down  192.0.2.125/24     node0         e0f     false
               data007    down/down  192.0.2.126/24     node0         e2a     false
               data008    down/down  192.0.2.127/24     node0         e2b     false
   8 entries were displayed.
   ```

If any LIFs appear with a `Status Admin` status of down or with an `Is home` status of `false`, continue with the next step.

2. Enable the data LIFs:

   `network interface modify {-role data} -status-admin up`

   Example

   ```
   cluster1::> network interface modify {-role data} -status-admin up
   8 entries were modified.
   ```

3. Revert LIFs to their home ports:

   `network interface revert *`

   Example

   This command reverts all LIFs back to their home ports and changes all LIF home statuses to `true`.

   ```
   cluster1::> network interface revert *
   8 entries were acted on.
   ```

4. Verify that all LIFs are in their home ports:

   `network interface show`

   Example

   This example shows that all LIFs for SVM vs0 are on their home ports.

   ```
   cluster1::> network interface show -vserver vs0
               Logical    Status     Network            Current       Current Is
   Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
   ----------- ---------- ---------- ------------------ ------------- ------- ----
   vs0
               data001      up/up    192.0.2.120/24     node0         e0e     true
               data002      up/up    192.0.2.121/24     node0         e0f     true
               data003      up/up    192.0.2.122/24     node0         e2a     true
               data004      up/up    192.0.2.123/24     node0         e2b     true
               data005      up/up    192.0.2.124/24     node1         e0e     true
               data006      up/up    192.0.2.125/24     node1         e0f     true
               data007      up/up    192.0.2.126/24     node1         e2a     true
               data008      up/up    192.0.2.127/24     node1         e2b     true
   8 entries were displayed.
   ```

Verifying the LIF failover configuration

Before you perform an upgrade, you use the `network interface show` command to verify the LIF failover configuration.

About this task

For more information about configuring and enabling LIF failover, see the *Clustered Data ONTAP Network Management Guide*.

Steps

1. Display the failover policy for each data port:
`network interface show -role data -failover`

Example

This example shows the default failover configuration for a two-node cluster with two data LIFs running Data ONTAP 8.1.x:

```
cluster1::> network interface show -role data -failover
         Logical         Home                  Failover        Failover
Vserver  Interface       Node:Port             Group Usage     Group
-------- --------------- --------------------- --------------- ---------------
vs0
         lif0            node0:e0b             system-defined
                             Failover Targets: node0:e0b, node0:e0c,
                                               node0:e0d, node0:e0e,
                                               node0:e0f, node1:e0b,
                                               node1:e0c, node1:e0d,
                                               node1:e0e, node1:e0f
vs1
         lif1            node1:e0b             system-defined
                             Failover Targets: node1:e0b, node1:e0c,
                                               node1:e0d, node1:e0e,
                                               node1:e0f, node0:e0b,
                                               node0:e0c, node0:e0d,
                                               node0:e0e, node0:e0f
```

This example shows the default failover configuration for a two-node cluster with two data LIFs running Data ONTAP 8.2.x:

```
cluster1::> network interface show -role data -failover
         Logical         Home                  Failover        Failover
Vserver  Interface       Node:Port             Policy          Group
-------- --------------- --------------------- --------------- ---------------
vs0
         lif0            node0:e0b             nextavail       system-defined
                         Failover Targets: node0:e0b, node0:e0c,
                                           node0:e0d, node0:e0e,
                                           node0:e0f, node1:e0b,
                                           node1:e0c, node1:e0d,
                                           node1:e0e, node1:e0f
vs1
```

```
          lif1            node1:e0b            nextavail       system-defined
                          Failover Targets: node1:e0b, node1:e0c,
                                            node1:e0d, node1:e0e,
                                            node1:e0f, node0:e0b,
                                            node0:e0c, node0:e0d,
                                            node0:e0e, node0:e0f
```

The `Failover Targets` field shows a prioritized list of failover targets for each LIF. For example, if lif0 fails over from its home port (e0b on node0), it will first attempt to fail over to port e0c on node0. If lif0 cannot fail over to e0c, it will next attempt to fail over to port e0d on node0, and so on.

2. If the data ports are on different VLANs or subnets, verify for each LIF that the failover group is user-defined.

   A user-defined failover group must be configured for each VLAN or broadcast domain, and each LIF must subscribe to the corresponding failover group. For more information about configuring LIF failover groups, see the *Clustered Data ONTAP Network Management Guide*.

3. If the failover policy is set to `disabled` for any of the LIFs, use the `network interface modify` command to enable failover.

4. For each LIF, verify that the `Failover Targets` field includes data ports from a different node that will remain up while the LIF's home node is being upgraded.

   You can use the `network interface failover-groups create` command to add a failover target to the failover group.

Disabling automatic LIF rebalancing

By disabling automatic LIF rebalancing before performing a batch upgrade, you can ensure that the LIFs remain online during the entire upgrade procedure.

About this task

When automatic LIF rebalancing is enabled, LIFs can be migrated to a less-utilized port on another node based on the LIF failover configuration. However, because a batch upgrade enables you to upgrade multiple nodes concurrently, automatic LIF rebalancing could cause the LIFs to migrate to a node that is rebooting.

Steps

1. Set the privilege level to advanced:

   `set -privilege advanced`

2. View and record any LIFs that have automatic LIF rebalancing enabled:

   `network interface show -allow-lb-migrate true`

Example

```
cluster1::*> network interface show -allow-lb-migrate true
            Logical    Status     Network            Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
----------- ---------- ---------- ------------------ ------------- ------- ----
vs0
            data1      up/up      192.0.2.120/24     node0         e0e     true
            data2      up/up      192.0.2.121/24     node0         e0f     true
            data3      up/up      192.0.2.122/24     node1         e0e     true
            data4      up/up      192.0.2.123/24     node1         e0f     true
            data5      up/up      192.0.2.124/24     node2         e0e     true
            data6      up/up      192.0.2.125/24     node2         e0f     true
            data7      up/up      192.0.2.126/24     node3         e0e     true
            data8      up/up      192.0.2.127/24     node3         e0f     true
8 entries were displayed.
```

You should record which LIFs have automatic rebalancing enabled so that you can reenable it after the batch upgrade is completed.

3. Disable automatic LIF rebalancing for each LIF that you identified:

   `network interface modify * -allow-lb-migrate false`

4. Return to the admin privilege level:

   `set -privilege admin`

Identifying active CIFS sessions that should be terminated

Before performing a minor nondisruptive upgrade or downgrade within the Data ONTAP 8.2 release family, you should identify and gracefully terminate any CIFS sessions that are not continuously available.

About this task

Continuously available CIFS shares, which are accessed by Hyper-V clients using the SMB3 protocol, do not need to be terminated before upgrading or downgrading.

Steps

1. Identify any established CIFS sessions that are not continuously available:

   `vserver cifs session show -continuously-available !Yes -instance`

   This command displays detailed information about any CIFS sessions that have no continuous availability.

Example

```
cluster1::> vserver cifs session show -continuously-available !Yes -
instance

                        Node: node1
```

```
                     Vserver: vs1
                  Session ID: 1
               Connection ID: 4160072788
Incoming Data LIF IP Address: 198.51.100.5
      Workstation IP address: 203.0.113.20
    Authentication Mechanism: NTLMv2
                Windows User: CIFSLAB\user1
                   UNIX User: nobody
                 Open Shares: 1
                  Open Files: 2
                  Open Other: 0
              Connected Time: 8m 39s
                   Idle Time: 7m 45s
            Protocol Version: SMB2_1
      Continuously Available: No
1 entry was displayed.
```

Each of the sessions identified by this command should be terminated before proceeding with the Data ONTAP upgrade or downgrade.

2. If necessary, identify the files that are open for each CIFS session that you identified:

   `vserver cifs session file show -session-id` ***`session_ID`***

   Example

   ```
   cluster1::> vserver cifs session file show -session-id 1

   Node:       node1
   Vserver:    vs1
   Connection: 4160072788
   Session:    1
   File    File      Open Hosting
   Continuously
   ID      Type      Mode Volume          Share                 Available
   ------- --------- ---- --------------- ---------------------
   ------------
   1       Regular   rw   vol10           homedirshare          No
   Path: \TestDocument.docx
   2       Regular   rw   vol10           homedirshare          No
   Path: \file1.txt
   2 entries were displayed.
   ```

Related concepts

Considerations for session-oriented protocols on page 163

Verifying the system time

You should verify that NTP is configured, and that the time is synchronized across the cluster.

About this task

For more information about managing the system time, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

Steps

1. Use the `system services ntp server show` command to verify that each node is associated with an NTP server.

   Example

   ```
   cluster1::> system services ntp server show
   Node   Server               Version
   ------ -------------------- -----------
   node0
          ntp1.example.com     max
          ntp2.example.com     max
   node1
          ntp1.example.com     max
          ntp2.example.com     max
   node2
          ntp1.example.com     max
          ntp2.example.com     max
   node3
          ntp1.example.com     max
          ntp2.example.com     max
   ```

2. Verify that each node has the same date and time:

If you are running Data ONTAP...	Enter the following command...
8.1.x	**`system node date show`**
8.2.x	**`cluster date show`**

   Example

   ```
   cluster1::> cluster date show
   Node      Date                Timezone
   --------- ------------------- -------------------------
   node0     4/6/2013 20:54:38   GMT
   node1     4/6/2013 20:54:38   GMT
   node2     4/6/2013 20:54:38   GMT
   node3     4/6/2013 20:54:38   GMT
   4 entries were displayed.
   ```

Determining the current software version on each node

For a successful upgrade, each node must be running the minimum software version allowed for the upgrade. You can check the software version by running the `system node image show` command.

About this task

Only clusters running Data ONTAP 8.1 or later can be upgraded directly to Data ONTAP 8.2.x.

Step

1. Determine the current software version:
   `system node image show`

   #### Example

```
cluster1::> system node image show
                 Is      Is                Install
Node     Image   Default Current Version    Date
-------- ------- ------- ------- --------- -------------------
node0
         image1  true    true    8.1.2     10/25/2012 12:36:46
         image2  false   false   8.1.0     11/27/2011 12:58:24
node1
         image1  true    true    8.1.2     10/25/2012 12:34:22
         image2  false   false   8.1.0     11/27/2011 12:58:26
4 entries were displayed.
```

Obtaining Data ONTAP software images

You must copy a software image from the N series support website (accessed and navigated as described in *Websites* on page 7) to an HTTP server on your network so that nodes can access the images by using the `system node image update` command.

About this task

To upgrade, revert, or downgrade the cluster to the target release of Data ONTAP, you need access to software images. Software images, firmware version information, and the latest firmware for your platform model are available on the N series support website (accessed and navigated as described in *Websites* on page 7). Note the following important information:

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

Steps

1. Locate the target Data ONTAP software on the N series support website (accessed and navigated as described in *Websites* on page 7).

2. Copy the software image (for example, `821_q_image.tgz`) from the N series support website (accessed and navigated as described in *Websites* on page 7) to the directory on the HTTP server from which the image will be served.

Installing Data ONTAP software images in a cluster

If needed, install the target Data ONTAP 8.x image software package, but leave the default set to the current Data ONTAP 8.x version.

Before you begin

You must have obtained the Data ONTAP software images.

Steps

1. Choose one of the following options based on your requirements:

If you want to...	Enter this command...
Download, but not install, the software image	**`system node image get -node * -package`** ***`location`*** **`-replace-package true -background true`** This command downloads the software image to all of the nodes simultaneously. To download the image to each node one at a time, do not specify the `-background` parameter.
Install a previously downloaded software image	**`system node image update -node * -package file:///mroot/etc/software/`*****`image_name`*** **`-background true`** Note the following considerations for this command: • If you are currently running Data ONTAP 8.2, the `-package` parameter only requires you to enter the image name; you do not need to enter the full directory path. • If you are unsure of the image name to install, you can view a list of previously downloaded software images by using the `system node image package show` command. • This command installs the software image on all of the nodes simultaneously. To install the image on each node one at a time, do not specify the `-background` parameter.
Download and install the software image in the same operation	**`system node image update -node * -package`** ***`location`*** **`-replace-package true -background true`** This command downloads and installs the software image on all of the nodes simultaneously. To download and install the image on each node one at a time, do not specify the `-background` parameter.

2. Verify that the software image is downloaded and installed on each node:

```
system node image show-update-progress -node *
```

This command displays the current status of the software image download and installation. You should continue to run this command until all nodes report a `Run Status` of `Exited`, and an `Exit Status` of `Success`.

Example

This example shows a 2-node cluster in which the software image has been downloaded and installed successfully on both nodes:

```
cluster1::> system node image show-update-progress -node *
There is no update/install in progress
Status of most recent operation:
        Run Status:     Exited
        Exit Status:    Success
        Phase:          Run Script
        Exit Message:   Installation complete. image2 updated on node node0.
There is no update/install in progress
Status of most recent operation:
        Run Status:     Exited
        Exit Status:    Success
        Phase:          Run Script
        Exit Message:   Installation complete. image2 updated on node node1.
2 entries were acted on.
```

How Data ONTAP software images are stored and alternated in the cluster

Each node in the cluster can hold two Data ONTAP software images, the current image that is running, and an alternate image that you can boot.

You can view the software images on each node in the cluster by using the `system node image show` command. The following example shows how the images alternate when you download a new software image and then upgrade to it.

How Data ONTAP images alternate

This example shows a two-node cluster. Each node has two images, *image1* (version 8.0.2) and *image2* (version 8.0.1). Both nodes are currently running *image1*, as indicated by the Is Current column. If you were to reboot either node, it would boot *image1*, because that is the default image (indicated by the Is Default column). The alternate image on both nodes is *image2*.

```
cluster1::> system node image show
                 Is      Is                Install
Node     Image   Default Current Version    Date
-------- ------- ------- ------- --------- -------------------
node0
         image1  true    true    8.0.2     3/25/2011 12:36:46
         image2  false   false   8.0.1     10/15/2010 12:58:24
node1
         image1  true    true    8.0.2     3/25/2011 12:34:22
         image2  false   false   8.0.1     10/15/2010 12:58:26
4 entries were displayed.
```

If you were to download a new software image, the new package would replace the package in the alternate image (*image2*) while all of the nodes continue to run the current image (*image1*).

```
cluster1::> system node image update -node * -package http://
10.28.99.99/8.1.1_q_image.tgz -replace-package true
```

After the download completes, the 8.1.1 version is now stored as the alternate image (*image2*), while both nodes continue to run version 8.0.2 (*image1*).

```
cluster1::> system node image show
                 Is      Is              Install
Node     Image   Default Current Version    Date
-------- ------- ------- ------- --------- -------------------
node0
         image1  true    true    8.0.2     3/25/2011 12:36:46
         image2  false   false   8.1.1     6/15/2012 10:23:18
node1
         image1  true    true    8.0.2     3/25/2011 12:34:22
         image2  false   false   8.1.1     6/15/2012 10:23:18
4 entries were displayed.
```

If you use the `system node image modify` command to set *image2* as the default image on the nodes, then after you reboot a node, it will boot *image2*. In this example, *image2* has been set to the default image, and the nodes have been rebooted. The images have alternated: *image1* is now the alternate image, and *image2* is the current image.

```
cluster1::> system node image show
                 Is      Is              Install
Node     Image   Default Current Version    Date
-------- ------- ------- ------- --------- -------------------
node0
         image1  false   false    8.0.2     3/25/2011 12:36:46
         image2  true    true     8.1.1     6/15/2012 10:23:18
node1
         image1  false   false    8.0.2     3/25/2011 12:34:22
         image2  true    true     8.1.1     6/15/2012 10:23:18
4 entries were displayed.
```

Preparing SnapMirror relationships for a nondisruptive upgrade or downgrade

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

About this task

For more information about SnapMirror operations, see the SnapMirror man pages and the *Clustered Data ONTAP Data Protection Guide*.

Steps

1. Use the `snapmirror show` command to determine the destination path for each SnapMirror relationship.
2. For each destination volume, suspend future SnapMirror transfers:

   `snapmirror quiesce -destination-path` ***`destination`***

   If there are no active transfers for the SnapMirror relationship, this command sets its status to `Quiesced`. If the relationship has active transfers, the status is set to `Quiescing` until the transfer is completed, and then the status becomes `Quiesced`.

   #### Example

   If you are upgrading from Data ONTAP 8.1, this example quiesces transfers involving the destination volume `vol1` from Storage Virtual Machine (SVM) `vs0` and cluster `cluster1`:

   ```
   cluster1::> snapmirror quiesce -destination-path cluster1://vs0/vol1
   ```

   #### Example

   If you are downgrading within the Data ONTAP 8.2 release family, this example quiesces transfers involving the destination volume `vol1` from SVM `vs0`:

   ```
   cluster1::> snapmirror quiesce -destination-path vs0:vol1
   ```

3. Verify that all SnapMirror relationships are quiesced:

   `snapmirror show -status !Quiesced`

   This command displays any SnapMirror relationships that are *not* quiesced.

   #### Example

   This example shows that all SnapMirror relationships are quiesced:

   ```
   cluster1::> snapmirror show -status !Quiesced
   There are no entries matching your query.
   ```

4. If any SnapMirror relationships are currently transferring, do one of the following options:

Option	Description
Wait for the transfers to complete before performing the Data ONTAP upgrade.	Once each transfer completes, the relationship changes to `Quiesced` status.

Option	Description
Stop the transfers by entering the following command: **`snapmirror abort -destination-path`** ***`destination`*** **`-h`** **Note:** You must use the `-foreground true` parameter if you are aborting load-sharing mirror transfers.	This command stops the SnapMirror transfer and restores the destination volume to the last Snapshot copy that was successfully transferred. The relationship is set to `Quiesced` status.

Ensuring that no jobs are running

You must verify the status of cluster jobs before upgrading or downgrading to a different Data ONTAP release. If any aggregate, volume, NDMP (dump or restore), or Snapshot jobs (such as create, delete, move, modify, replicate, and mount jobs) are running or queued, allow the jobs to finish successfully or stop the queued entries.

Steps

1. Review the list of any running or queued aggregate, volume, or Snapshot jobs:

   `job show`

   #### Example

   ```
   cluster1::> job show
                               Owning
   Job ID Name                 Vserver    Node           State
   ------ -------------------- ---------- -------------- ----------
   8629   Vol Reaper           cluster1   -              Queued
          Description: Vol Reaper Job
   8630   Certificate Expiry Check
                               cluster1   -              Queued
          Description: Certificate Expiry Check
   8632   CLUSTER BACKUP AUTO daily
                               cluster1   -              Queued
          Description: Cluster Backup Job
   8633   CLUSTER BACKUP AUTO weekly
                               cluster1   -              Queued
          Description: Cluster Backup Job
   9944   SnapMirrorDaemon_7_2147484678
                               cluster1   node1          Dormant
          Description: Snapmirror Daemon for 7_2147484678
   18277  CLUSTER BACKUP AUTO 8hour
                               cluster1   -              Queued
          Description: Cluster Backup Job
   18377  SnapMirror Service Job
                               cluster1   node0          Dormant
          Description: SnapMirror Service Job
   18379  Network Consistency Diagnostic - weekly
                               cluster1   node0          Queued
          Description: Network Consistency Checker
   18385  Network Consistency Diagnostic - weekly
                               cluster1    node1         Queued
          Description: Network Consistency Checker
   9 entries were displayed
   ```

2. Delete any running or queued aggregate, volume, or Snapshot copy jobs:

job delete -id *job_id*

Example

```
cluster1::> job delete -id 8629
```

3. Ensure that no aggregate, volume, or Snapshot jobs are running or queued:

   job show

Example

In this example, all running and queued jobs have been deleted.

```
cluster1::> job show
                            Owning
Job ID Name                 Vserver    Node           State
------ -------------------- ---------- -------------- ----------
9944   SnapMirrorDaemon_7_2147484678
                            cluster1   node1          Dormant
       Description: Snapmirror Daemon for 7_2147484678
18377  SnapMirror Service Job
                            cluster1   node0          Dormant
       Description: SnapMirror Service Job
2 entries were displayed
```

Performing the software upgrade

To upgrade a cluster to a new Data ONTAP release, you must verify that the cluster is ready to be upgraded, choose an upgrade method, and then perform the steps for the upgrade method.

You can select one of the following upgrade methods:

- Rolling upgrade (nondisruptive)
- Batch upgrade (nondisruptive)
- Disruptive upgrade

Related concepts

Types of cluster upgrades on page 11

Verifying that the cluster is ready to be upgraded

You must verify that the target Data ONTAP software is installed, storage failover is enabled, and if necessary, cluster HA is enabled.

Steps

1. Verify that the Data ONTAP 8.2 software is installed:

   system node image show

Example

This example shows that version 8.2.1 is installed as the alternate image on both nodes.

```
cluster1::> system node image show
                 Is      Is                 Install
Node     Image   Default Current Version    Date
-------- ------- ------- ------- --------   -------------------
node0
         image1  true    true    8.1.2      10/25/2012 12:37:36
         image2  false   false   8.2.1      11/22/2013 13:52:22
node1
         image1  true    true    8.1.2      10/25/2012 12:41:16
         image2  false   false   8.2.1      11/22/2013 13:55:22
4 entries were displayed.
```

For more information about installing the target Data ONTAP software image, see *Installing Data ONTAP 8.x software images in a cluster*.

2. Set the Data ONTAP 8.2 software image to be the default image:

   `system image modify {-node * -iscurrent false} -isdefault true`

3. Verify that the Data ONTAP 8.2 software image is set as the default image:

   `system node image show`

Example

This example shows that version 8.2.1 is set as the default image on both nodes.

```
cluster1::> system node image show
                 Is      Is                 Install
Node     Image   Default Current Version    Date
-------- ------- ------- ------- --------   -------------------
node0
         image1  false    true   8.1.2      10/25/2012 12:37:36
         image2  true     false  8.2.1      11/22/2013 13:52:22
node1
         image1  false   true    8.1.2      10/25/2012 12:41:16
         image2  true    false   8.2.1      11/22/2013 13:55:22
4 entries were displayed.
```

4. If you are performing a nondisruptive upgrade, verify the high availability configuration:

   a) Verify that storage failover is enabled and possible:

   `storage failover show`

Example

This example shows that storage failover is enabled and possible on node0 and node1:

```
cluster1::> storage failover show
                              Takeover
Node           Partner        Possible State
-------------- -------------- -------- -------------------------------------
node0          node1          true     Connected to node1
node1          node0          true     Connected to node0
2 entries were displayed.
```

You can enable storage failover by using the `storage failover modify` command.

b) If the cluster consists of only two nodes (a single HA pair), ensure that cluster HA is configured:

`cluster ha show`

Example

```
cluster1::> cluster ha show
High Availability Configured: true
```

You can enable cluster HA by using the `cluster ha modify` command.

Upgrading a Data ONTAP cluster nondisruptively by using the rolling upgrade method

This nondisruptive upgrade (NDU) method has several steps: initiating a failover operation on each node in an HA pair, updating the "failed" node, initiating giveback, and then repeating the process for each HA pair in the cluster.

Before you begin

You must have satisfied upgrade preparation requirements and verified that the cluster is ready to be upgraded.

Steps

1. *Upgrading the first node in an HA pair* on page 49

You upgrade the first node in an HA pair by initiating a takeover by the node's partner. The partner serves the node's data while the first node is upgraded.

2. *Upgrading the second node in an HA pair* on page 54

After upgrading the first node in an HA pair, you upgrade its partner by initiating a takeover on it. The first node serves the partner's data while the partner node is upgraded.

3. *Verifying that the HA pair was upgraded successfully* on page 58

After upgrading both nodes in an HA pair, you must verify that the target release is running on both nodes.

4. Repeat Steps 1-3 for each additional HA pair.

After you finish

You should complete post-upgrade tasks.

Related references

Cluster upgrade checklist on page 11

Upgrading the first node in an HA pair

You upgrade the first node in an HA pair by initiating a takeover by the node's partner. The partner serves the node's data while the first node is upgraded.

Steps

1. Disable automatic giveback on both nodes of the HA pair if it is enabled by entering the following command on each node:

   `storage failover modify -node` ***`nodename`*** **`-auto-giveback false`**

   If the cluster is a two-node cluster, a message is displayed warning you that disabling automatic giveback will prevent the management cluster services from going online in the event of a alternating-failure scenario. Enter **`y`** to continue.

2. Verify that automatic giveback is disabled for both nodes:

   `storage failover show -fields auto-giveback`

   #### Example

   ```
   cluster1::> storage failover show -fields auto-giveback
   node     auto-giveback
   -------- -------------
   node0    false
   node1    false
   2 entries were displayed.
   ```

3. Determine if the node to be upgraded is currently serving any clients by entering the following command twice:

   `system node run -node` ***`nodenameA`*** **`-command uptime`**

   The `uptime` command displays the total number of operations the node has performed for NFS, CIFS, FC, and iSCSI clients since the node was last booted. For each protocol, determine if the operation counts are increasing. If they are increasing, the node is currently serving clients for that protocol. If they are not increasing, the node is not currently serving clients for that protocol.

   You should make a note of each protocol that has increasing client operations so that after the node is upgraded, you can verify that client traffic has resumed.

   #### Example

   This example shows a node with NFS, CIFS, FC, and iSCSI operations. However, the node is currently serving only NFS and iSCSI clients.

   ```
   cluster1::> system node run -node node0 -command uptime
     2:58pm up  7 days, 19:16 800000260 NFS ops, 1017333 CIFS ops, 0 HTTP ops, 40395 FCP ops, 32810 iSCSI
   ops
   ```

```
cluster1::> system node run -node node0 -command uptime
  2:58pm up  7 days, 19:17 800001573 NFS ops, 1017333 CIFS ops, 0 HTTP ops, 40395 FCP ops, 32815 iSCSI
ops
```

4. Migrate LIFs away from the node:

```
network interface migrate-all -node nodenameA
```

Data LIFs for SAN protocols are not migrated. As long as these LIFs exist on each node in the cluster, data can be served through alternate paths during the upgrade process.

If you are connected to the cluster through the cluster management LIF, and if this node currently hosts the cluster management LIF, then your SSH session will be temporarily disconnected while the LIF is migrated. After the LIF has migrated, you can log in to the cluster through the cluster management LIF.

5. Verify that the LIFs migrated away from the node:

```
network interface show -data-protocol nfs,cifs -role data -curr-node
nodenameA
```

Example

This example shows that node0's data LIFs were migrated to a different node.

```
cluster1::> network interface show -data-protocol nfs,cifs -role data -curr-node node0
 There are no entries matching your query.
```

6. Trigger an AutoSupport notification:

```
system node autosupport invoke -node nodenameA -type all -message
"starting_NDU"
```

This AutoSupport notification includes a record of the system status just prior to upgrade. It saves useful troubleshooting information in case there is a problem with the upgrade process.

If your cluster is not configured to send AutoSupport messages, a copy of the notification is saved locally.

7. Initiate a takeover:

```
storage failover takeover -ofnode nodenameA
```

Do not specify the parameter `-option immediate`, because a normal takeover is required for the node that is being taken over to boot onto the new software image.

The first node boots up to the `Waiting for giveback` state.

> **Note:** If AutoSupport is enabled, an AutoSupport message is sent indicating that the node is out of cluster quorum. You can safely ignore this notification and proceed with the upgrade.

8. Verify that the takeover was successful:

```
storage failover show
```

Example

This example shows that the takeover was successful. Node node0 is in the `Waiting for giveback` state, and its partner is `In takeover`.

```
cluster1::> storage failover show
                              Takeover
Node           Partner        Possible State Description
-------------- -------------- -------- -------------------------------------
node0          node1          -        Waiting for giveback (HA mailboxes)
node1          node0          false    In takeover
2 entries were displayed.
```

9. Wait 8 minutes to ensure the following conditions:

   - Client multipathing (if deployed) is stabilized.
   - Clients are recovered from the pause in I/O that occurs during takeover.

     The recovery time is client-specific and may take longer than 8 minutes depending on the characteristics of the client applications.

10. Return the aggregates to the first node:

    `storage failover giveback –ofnode` ***`nodenameA`***

    Attention: The giveback is not initiated, an error message is returned, and an event is generated if any conditions such as the following are detected:

    - Long-running operations (such as ASUP generation)
    - Operations that cannot be restarted (such as aggregate creation)
    - Error conditions (such as a disk connectivity mismatch between the nodes)

    If giveback is not initiated, complete the following steps:

    a. Address the "veto" condition described in the error message, ensuring that any identified operations are terminated gracefully.

    b. Reenter the giveback command:

       `storage failover giveback -ofnode` ***`nodenameA`***

    Alternatively, you can analyze the messages and events for relevance in your environment. If you determine that the veto conditions are not significant, you can override the giveback veto by entering the following command:

    `storage failover giveback –ofnode` ***`nodenameA`*** **`-override-vetoes true`**

    For more information about determining whether you can safely override the veto, see the *Clustered Data ONTAP High-Availability Configuration Guide*.

    This first returns the root aggregate to the partner node and then, after that node has finished booting, returns the non-root aggregates.

11. Verify that all aggregates have been returned:

    `storage failover show-giveback`

If the `Giveback Status` field indicates that the node that was taken over is in partial giveback, then complete the following steps before proceeding:

a) Determine which aggregates have been returned:

   `storage aggregate show -node` ***`nodenameA`***

b) Check EMS logs for errors and take corrective action.

   For more information about EMS messages, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

c) Repeat Step 10.a to verify any corrections.

The newly-booted node begins to serve data to clients from each aggregate as soon as the aggregate is returned.

12. Verify that the upgrade completed successfully for the node:

a) Set the privilege level to advanced:

   `set -privilege advanced`

b) Ensure that upgrade status is complete for the node:

   `system node upgrade-revert show -node` ***`nodenameA`***

   The status should be listed as `complete`.

   If the status is not successful, from the node, run the `system node upgrade-revert upgrade` command. If this command does not complete the node's upgrade, contact technical support immediately.

c) Return to the admin privilege level:

   `set -privilege admin`

13. Revert the LIFs back to the node:

`network interface revert *`

Example

This command returns the LIFs that were migrated away from the node.

```
cluster1::> network interface revert *
8 entries were acted on.
```

14. Verify that the node's data ports and LIFs are up and operational:

a) Verify that the node's data ports are up:

   `network port show -node` ***`nodenameA`*** **`-role data`**

Example

This example shows that all of the node's data ports are up.

```
cluster1::> network port show -node node0 -role data
                                      Auto-Negot  Duplex     Speed (Mbps)
```

```
Node   Port   Role         Link   MTU Admin/Oper  Admin/Oper Admin/Oper
------ ------ ------------ ---- ----- ----------- ---------- ------------
node0
       e0c    data         up    9000  true/true  full/full   auto/1000
       e0d    data         up    9000  true/true  full/full   auto/1000
       e1b    data         up    9000  true/true  full/full   auto/1000
       e1c    data         up    9000  true/true  full/half   auto/10
       e1d    data         up    9000  true/true  full/half   auto/10
5 entries were displayed.
```

b) Verify that the node's data LIFs successfully reverted back to the node, and that they are up:

network interface show -data-protocol nfs,cifs -role data -curr-node *nodenameA*

Example

This example shows that all of the data LIFs hosted by the node have successfully reverted back to the node, and that they are operationally up.

```
cluster1::> network interface show -data-protocol nfs,cifs -role data -curr-node node0
            Logical    Status     Network            Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
----------- ---------- ---------- ------------------ ------------- ------- ----
vs0
            data001      up/up    192.0.2.120/24     node0         e0c     true
            data002      up/up    192.0.2.121/24     node0         e0d     true
            data003      up/up    192.0.2.122/24     node0         e0d     true
            data004      up/up    192.0.2.123/24     node0         e0c     true
4 entries were displayed.
```

15. If you previously determined that this node serves clients, verify that the node is providing service for each protocol that it was previously serving:

system node run -node *nodenameA* -command uptime

The operation counts reset to zero during the upgrade.

Example

This example shows that the upgraded node has resumed serving its NFS and iSCSI clients:

```
cluster1::> system node run -node node0 -command uptime
  3:15pm up  0 days, 0:16 129 NFS ops, 0 CIFS ops, 0 HTTP ops, 0 FCP ops, 2 iSCSI ops
```

16. Trigger an AutoSupport notification:

system node autosupport invoke -node *nodenameA* -type all -message "finishing_NDU"

Upgrading the partner node in an HA pair

After upgrading the first node in an HA pair, you upgrade its partner by initiating a takeover on it. The first node serves the partner's data while the partner node is upgraded.

Steps

1. Determine if the node to be upgraded is currently serving any clients by entering the following command twice:

   `system node run -node` ***`nodenameB`*** **`-command uptime`**

   The `uptime` command displays the total number of operations the node has performed for NFS, CIFS, FC, and iSCSI clients since the node was last booted. For each protocol, determine if the operation counts are increasing. If they are increasing, the node is currently serving clients for that protocol. If they are not increasing, the node is not currently serving clients for that protocol.

   You should make a note of each protocol that has increasing client operations so that after the node is upgraded, you can verify that client traffic has resumed.

   #### Example

   This example shows a node with NFS, CIFS, FC, and iSCSI operations. However, the node is currently serving only NFS and iSCSI clients.

   ```
   cluster1::> system node run -node node1 -command uptime
     2:58pm up  7 days, 19:16 800000260 NFS ops, 1017333 CIFS ops, 0 HTTP ops, 40395 FCP ops, 32810 iSCSI
   ops

   cluster1::> system node run -node node1 -command uptime
     2:58pm up  7 days, 19:17 800001573 NFS ops, 1017333 CIFS ops, 0 HTTP ops, 40395 FCP ops, 32815 iSCSI
   ops
   ```

2. Migrate LIFs away from the node:

   `network interface migrate-all -node` ***`nodenameB`***

   Data LIFs for SAN protocols are not migrated. As long as these LIFs exist on each node in the cluster, data can be served through alternate paths during the upgrade process.

   If you are connected to the cluster through the cluster management LIF, and if this node currently hosts the cluster management LIF, then your SSH session will be temporarily disconnected while the LIF is migrated. After the LIF has migrated, you can log in to the cluster through the cluster management LIF.

3. Verify that the LIFs migrated to their failover targets:

   `network interface show -data-protocol nfs,cifs -role data -curr-node` ***`nodenameB`***

   #### Example

   This example shows that node1's data LIFs were migrated to a different node.

```
cluster1::> network interface show -data-protocol nfs,cifs -role data -curr-node node1
 There are no entries matching your query.
```

4. Trigger an AutoSupport notification:

   system node autosupport invoke -node *nodenameB* -type all -message "starting_NDU"

   This AutoSupport notification includes a record of the system status just prior to upgrade. It saves useful troubleshooting information in case there is a problem with the upgrade process.

   If your cluster is not configured to send AutoSupport messages, a copy of the notification is saved locally.

5. Initiate a takeover by using one of the following commands:

If you are upgrading from a...	Enter this command...
Data ONTAP 8.2.x release	**storage failover takeover -ofnode *nodenameB***
Data ONTAP 8.1.x release	**storage failover takeover -ofnode *nodenameB* -option allow-version-mismatch** The `allow-version-mismatch` option enables the HA pair to tolerate different Data ONTAP release family versions during a major release upgrade.

   Do not specify the parameter `-option immediate`, because a normal takeover is required for the node that is being taken over to boot onto the new software image.

   The node that is taken over boots up to the `Waiting for giveback` state.

   Note: If AutoSupport is enabled, an AutoSupport message is sent indicating that the node is out of cluster quorum. You can safely ignore this notification and proceed with the upgrade.

6. Verify that the takeover was successful:

   storage failover show

Example

This example shows that the takeover was successful. Node node1 is in the `Waiting for giveback` state, and its partner is `In takeover`.

```
cluster1::> storage failover show
                              Takeover
Node           Partner        Possible State Description
-------------- -------------- -------- -------------------------------------
node0          node1          -        In takeover
node1          node0          false    Waiting for giveback (HA mailboxes)
2 entries were displayed.
```

7. Wait 8 minutes to ensure the following conditions:

   - Client multipathing (if deployed) is stabilized.

- Clients are recovered from the pause in I/O that occurs during takeover.

  The recovery time is client-specific and may take longer than 8 minutes depending on the characteristics of the client applications.

8. Return the aggregates to the partner node:

```
storage failover giveback -ofnode nodenameB
```

Attention: The giveback is not initiated, an error message is returned, and an event is generated if any conditions such as the following are detected:

- Long-running operations (such as ASUP generation)
- Operations that cannot be restarted (such as aggregate creation)
- Error conditions (such as a disk connectivity mismatch between the nodes)

If giveback is not initiated, complete the following steps:

a. Address the "veto" condition described in the error message, ensuring that any identified operations are terminated gracefully.

b. Reenter the giveback command:

```
storage failover giveback -ofnode nodenameB
```

Alternatively, you can analyze the messages and events for relevance in your environment. If you determine that the veto conditions are not significant, you can override the giveback veto by entering the following command:

```
storage failover giveback -ofnode nodenameB -override-vetoes true
```

For more information about determining whether you can safely override the veto, see the *Clustered Data ONTAP High-Availability Configuration Guide*.

This first returns the root aggregate to the partner node and then, after that node has finished booting, returns the non-root aggregates.

9. Verify that all aggregates have been returned:

```
storage failover show-giveback
```

If the `Giveback Status` field indicates that the node that was taken over is in partial giveback, then complete the following steps before proceeding:

a) Determine which aggregates have been returned:

```
storage aggregate show -node nodenameB
```

b) Check EMS logs for errors and take corrective action.

For more information about EMS messages, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

c) Repeat Step 8.a to verify any corrections.

The newly-booted node begins to serve data to clients from each aggregate as soon as the aggregate is returned.

10. Verify that the upgrade completed successfully for the node:

a) Set the privilege level to advanced:

`set -privilege advanced`

b) Ensure that upgrade status is complete for the node:

`system node upgrade-revert show -node` ***`nodenameB`***

The status should be listed as `complete`.

If the status is not successful, from the node, run the `system node upgrade-revert upgrade` command. If this command does not complete the node's upgrade, contact technical support immediately.

c) Return to the admin privilege level:

`set -privilege admin`

11. Revert the LIFs back to the node:

`network interface revert *`

This command returns the LIFs that were migrated away from the node.

Example

```
cluster1::> network interface revert *
8 entries were acted on.
```

12. Verify that the node's data ports and LIFs are up and operational:

a) Verify that the node's data ports are up:

`network port show -node` ***`nodenameB`*** **`-role data`**

Example

This example shows that all of the node's data ports are up.

```
cluster1::> network port show -node node1 -role data
                                      Auto-Negot  Duplex     Speed (Mbps)
Node   Port   Role         Link   MTU Admin/Oper  Admin/Oper Admin/Oper
------ ------ ------------ ---- ----- ----------- ---------- ------------
node1
       e0c    data         up    9000  true/true  full/full   auto/1000
       e0d    data         up    9000  true/true  full/full   auto/1000
       e1b    data         up    9000  true/true  full/full   auto/1000
       e1c    data         up    9000  true/true  full/half   auto/10
       e1d    data         up    9000  true/true  full/half   auto/10
5 entries were displayed.
```

b) Verify that the node's data LIFs successfully reverted back to the node, and that they are up:

`network interface show -data-protocol nfs,cifs -role data -curr-node` ***`nodenameB`***

Example

This example shows that all of the data LIFs hosted by the node have successfully reverted back to the node, and that they are operationally up.

```
cluster1::> network interface show -data-protocol nfs,cifs -role data -curr-node node1
            Logical    Status     Network            Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
----------- ---------- ---------- ------------------ ------------- ------- ----
vs0
            data001      up/up    192.0.2.120/24     node1         e0c     true
            data002      up/up    192.0.2.121/24     node1         e0d     true
            data003      up/up    192.0.2.122/24     node1         e0d     true
            data004      up/up    192.0.2.123/24     node1         e0c     true
4 entries were displayed.
```

13. If you previously determined that this node serves clients, verify that the node is providing service for each protocol that it was previously serving:

    `system node run -node` *`nodenameB`* `-command uptime`

    The operation counts reset to zero during the upgrade.

    #### Example

    This example shows that the upgraded node has resumed serving its NFS and iSCSI clients:

    ```
    cluster1::> system node run -node node1 -command uptime
      3:15pm up  0 days, 0:16 129 NFS ops, 0 CIFS ops, 0 HTTP ops, 0 FCP ops, 2 iSCSI ops
    ```

14. Trigger an AutoSupport notification:

    `system node autosupport invoke -node` *`nodenameB`* `-type all -message "finishing_NDU"`

Verifying that the HA pair was upgraded successfully

After upgrading both nodes in an HA pair, you must verify that the target release is running on both nodes.

Steps

1. Confirm that the new Data ONTAP 8.2.x software is running on both nodes of the HA pair:

    `system node image show`

    #### Example

    This example shows version 8.2.1 as the current version on both nodes:

    ```
    cluster1::> system node image show
                      Is      Is                 Install
    Node     Image    Default Current Version    Date
    -------- -------  ------- ------- --------   -------------------
    node0
             image1   true    true    8.2.1      11/22/2013 13:52:22
    ```

```
         image2  false   false   8.1.2      10/25/2012 12:37:36
node1
         image1  true    true    8.2.1      11/22/2013 13:55:22
         image2  false   false   8.1.2      10/25/2012 12:41:16
4 entries were displayed.
```

2. Re-enable automatic giveback on both nodes if it was previously disabled:

   `storage failover modify -node` ***`nodename`*** **`-auto-giveback true`**

3. Ensure that the cluster is in quorum and that services are running before upgrading the next pair of nodes.

   You can use the `cluster show` and `cluster ring show` commands to verify that the cluster is in quorum.

After you finish

Upgrade any additional HA pairs. After all of the HA pairs are upgraded, you should verify that the cluster was upgraded successfully.

Upgrading a Data ONTAP cluster nondisruptively by using the batch method

If your cluster contains eight or more nodes, you can upgrade Data ONTAP by dividing the cluster into two upgrade batches, upgrading a set of nodes in the first batch, upgrading their high-availability partners, and then repeating the process for the second batch.

Before you begin

- You must have completed the upgrade preparation requirements and verified that the cluster is ready to be upgraded.
- You must have determined the upgrade sequence for the batch upgrade.

About this task

If the cluster serves SAN clients, do not use the batch upgrade method. You should perform a rolling upgrade instead.

Steps

1. *Upgrading the first set of nodes in a batch of HA pairs* on page 60

   You upgrade the first set of nodes in a batch of HA pairs by initiating a takeover by the nodes' partners. The partners serve the nodes' data while the first set of nodes is upgraded.

2. *Upgrading the second set of nodes in a batch of HA pairs* on page 65

   You upgrade the partner nodes in a batch of HA pairs by initiating a takeover on the nodes. The first set of nodes serve the nodes' data while the partners are upgraded.

3. *Verifying that the batch was upgraded successfully* on page 70

After upgrading all of the nodes in a batch, you must verify that the target release is running on the nodes.

4. Repeat Steps 1-3 to upgrade the second batch.

After you finish

You should verify that the cluster was upgraded successfully.

Related concepts

Required upgrade sequence for a batch upgrade on page 19

Related references

Cluster upgrade checklist on page 11

Upgrading the first set of nodes in a batch of HA pairs

You upgrade the first set of nodes in a batch of HA pairs by initiating a takeover by the nodes' partners. The partners serve the nodes' data while the first set of nodes is upgraded.

Steps

1. Disable automatic giveback on the HA pairs in the first batch by entering the following command on each node:

   `storage failover modify -node` ***`nodename`*** **`-auto-giveback false`**

2. Verify that automatic giveback is disabled for the nodes in the batch:

   `storage failover show -fields auto-giveback`

   #### Example

   This example shows that automatic giveback has been disabled on all of the nodes in the batch.

   ```
   cluster1::> storage failover show -fields auto-giveback
   node     auto-giveback
   -------- -------------
   node0    false
   node1    false
   node2    false
   node3    false
   node4    false
   node5    false
   node6    true
   node7    true
   node8    true
   node9    true
   node10   true
   node11   true
   12 entries were displayed.
   ```

3. Determine if the nodes to be upgraded are currently serving any clients by entering the following command twice for each node:

`system node run -node` *`nodename`* `-command uptime`

The `uptime` command displays the total number of operations the node has performed for NFS, CIFS, FC, and iSCSI clients since the node was last booted. For each protocol, determine if the operation counts are increasing. If they are increasing, the node is currently serving clients for that protocol. If they are not increasing, the node is not currently serving clients for that protocol.

You should make a note of each protocol that has increasing client operations so that after the node is upgraded, you can verify that client traffic has resumed.

Example

This example shows a node with NFS, CIFS, FC, and iSCSI operations. However, the node is currently serving only NFS and iSCSI clients.

```
cluster1::> system node run -node node0 -command uptime
  2:58pm up  7 days, 19:16 800000260 NFS ops, 1017333 CIFS ops, 0 HTTP ops, 40395 FCP ops, 32810 iSCSI
ops

cluster1::> system node run -node node0 -command uptime
  2:58pm up  7 days, 19:17 800001573 NFS ops, 1017333 CIFS ops, 0 HTTP ops, 40395 FCP ops, 32815 iSCSI
ops
```

4. Migrate LIFs away from the nodes that will be taken over by entering the following command for each node in the first set:

   `network interface migrate {-data-protocol nfs,cifs -role data -curr-node` *`source_node`*`} -dest-node` *`partner_node`*

   Data LIFs for SAN protocols are not migrated. As long as these LIFs exist on each node in the cluster, data can be served through alternate paths during the upgrade process.

5. Verify that the LIFs migrated to the proper ports on the nodes' partners by entering the following command for each node's partner:

   `network interface show -data-protocol nfs,cifs -role data -curr-node` *`partner_node`*

Example

This example shows that node0's data LIFs were migrated to port e0b on its partner (node1).

```
cluster1::> network interface show -data-protocol nfs,cifs -role data -curr-node node1
            Logical    Status     Network            Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
----------- ---------- ---------- ------------------ ------------- ------- ----
vs0
            lif1       up/up      192.0.2.130/24     node1         e0b     true
            lif2       up/up      192.0.2.131/24     node1         e0b     false
            lif3       up/up      192.0.2.132/24     node1         e0b     true
vs1
            lif1       up/up      192.0.2.133/24     node1         e0b     false
            lif2       up/up      192.0.2.134/24     node1         e0b     true
5 entries were displayed.
```

If desired, you can migrate a LIF to a different port on the partner node by using the `network interface migrate` command.

6. Trigger an AutoSupport notification by entering the following command for each node in the set:

   `system node autosupport invoke -node` ***`nodename`*** **`-type all -message "starting_NDU"`**

   This AutoSupport notification includes a record of the system status just prior to upgrade. It saves useful troubleshooting information in case there is a problem with the upgrade process.

   If your cluster is not configured to send AutoSupport messages, a copy of the notification is saved locally.

7. Initiate a takeover by entering the following command for each node in the first set:

   `storage failover takeover -ofnode` ***`nodename`***

   Do not specify the parameter `-option immediate`, because a normal takeover is required for the nodes that are being taken over to boot onto the new software image.

   The nodes boot up to the `Waiting for giveback` state.

   > **Note:** If AutoSupport is enabled, an AutoSupport message is sent indicating that the nodes are out of cluster quorum. You can safely ignore this notification and proceed with the upgrade.

8. Verify that the takeover was successful:

   `storage failover show`

   Example

   This example shows that the takeover was successful. The first set of nodes (node0, node2, and node4) are in the `Waiting for giveback` state, and their partners are `In takeover`.

```
cluster1::> storage failover show
                              Takeover
Node           Partner        Possible State Description
-------------- -------------- -------- -------------------------------------
node0          node1          -        Waiting for giveback (HA mailboxes)
node1          node0          false    In takeover
node2          node3          -        Waiting for giveback (HA mailboxes)
node3          node2          false    In takeover
node4          node5          -        Waiting for giveback (HA mailboxes)
node5          node4          false    In takeover
node6          node7          true     Connected to node7
node7          node6          true     Connected to node6
node8          node9          true     Connected to node9
node9          node8          true     Connected to node8
node10         node11         true     Connected to node11
node11         node10         true     Connected to node10
12 entries were displayed.
```

9. Wait 8 minutes to ensure the following conditions:

   - Client multipathing (if deployed) is stabilized.
   - Clients are recovered from the pause in I/O that occurs during takeover.

     The recovery time is client-specific and may take longer than 8 minutes depending on the characteristics of the client applications.

10. Return the aggregates to the nodes by entering the following command for each of the nodes' partners:

```
storage failover giveback –ofnode nodename
```

Attention: The giveback is not initiated, an error message is returned, and an event is generated if any conditions such as the following are detected:

- Long-running operations (such as ASUP generation)
- Operations that cannot be restarted (such as aggregate creation)
- Error conditions (such as a disk connectivity mismatch between the nodes)

If giveback is not initiated, complete the following steps:

a. Address the "veto" condition described in the error message, ensuring that any identified operations are terminated gracefully.

b. Reenter the giveback command:

```
storage failover giveback -ofnode nodename
```

Alternatively, you can analyze the messages and events for relevance in your environment. If you determine that the veto conditions are not significant, you can override the giveback veto by entering the following command:

```
storage failover giveback –ofnode nodename -override-vetoes true
```

For more information about determining whether you can safely override the veto, see the *Clustered Data ONTAP High-Availability Configuration Guide*.

This first returns the root aggregate to the partner nodes and then, after those nodes have finished booting, returns the non-root aggregates.

11. Verify that all aggregates have been returned:

```
storage failover show-giveback
```

If the `Giveback Status` field indicates that a node that was taken over is in partial giveback, then complete the following steps before proceeding:

a) Determine which aggregates have been returned:

```
storage aggregate show -node nodename
```

b) Check EMS logs for errors and take corrective action.

For more information about EMS messages, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

c) Repeat Step 10.a to verify any corrections.

The newly-booted nodes begin to serve data to clients from each aggregate as soon as the aggregate is returned.

12. Verify that the upgrade completed successfully for the nodes in the first set:

a) Set the privilege level to advanced:

```
set -privilege advanced
```

b) Ensure that upgrade status is complete for each node in the first set:

`system node upgrade-revert show -node` ***`nodename`***

The status should be listed as `complete`.

If the status is not successful for any node in the set, from the node, run the `system node upgrade-revert upgrade` command. If this command does not complete the node's upgrade, contact technical support immediately.

c) Return to the admin privilege level:

`set -privilege admin`

13. Revert the LIFs back to the first set of nodes:

`network interface revert *`

Example

This command returns the LIFs that were migrated away from the nodes.

```
cluster1::> network interface revert *
12 entries were acted on.
```

14. Verify that the nodes' data ports and LIFs are up and operational by completing the following steps for each node in the first set:

a) Verify that the data ports for are up for each node in the first set:

`network port show -node` ***`nodename`*** **`-role data`**

Example

This example shows that all of a node's data ports are up.

```
cluster1::> network port show -node node0 -role data
                                      Auto-Negot  Duplex     Speed (Mbps)
Node   Port   Role         Link   MTU Admin/Oper  Admin/Oper Admin/Oper
------ ------ ------------ ---- ----- ----------- ---------- ------------
node0
       e0c    data         up    9000  true/true  full/full   auto/1000
       e0d    data         up    9000  true/true  full/full   auto/1000
       e1b    data         up    9000  true/true  full/full   auto/1000
       e1c    data         up    9000  true/true  full/half   auto/10
       e1d    data         up    9000  true/true  full/half   auto/10
5 entries were displayed.
```

b) Verify that data LIFs successfully reverted back to each node in the first set, and that the LIFs are up:

`network interface show -data-protocol nfs,cifs -role data -curr-node` ***`nodename`***

Example

This example shows that all of the data LIFs hosted by a node have successfully reverted back to the node, and that they are operationally up.

```
cluster1::> network interface show -data-protocol nfs,cifs -role data -curr-node node0
            Logical    Status     Network            Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
----------- ---------- ---------- ------------------ ------------- ------- ----
vs0
            data001      up/up    192.0.2.120/24     node0         e0c     true
            data002      up/up    192.0.2.121/24     node0         e0d     true
            data003      up/up    192.0.2.122/24     node0         e0d     true
            data004      up/up    192.0.2.123/24     node0         e0c     true
4 entries were displayed.
```

15. If you previously determined that any of the nodes in the first set serve clients, verify that each node is providing service for each protocol that it was previously serving:

    `system node run -node` ***`nodename`*** **`-command uptime`**

    The operation counts reset to zero during the upgrade.

    #### Example

    This example shows that an upgraded node has resumed serving its NFS and iSCSI clients:

    ```
    cluster1::> system node run -node node0 -command uptime
      3:15pm up  0 days, 0:16 129 NFS ops, 0 CIFS ops, 0 HTTP ops, 0 FCP ops, 2 iSCSI ops
    ```

16. Trigger an AutoSupport notification by entering the following command for each of the nodes in the first set:

    `system node autosupport invoke -node` ***`nodename`*** **`-type all -message "finishing_NDU"`**

Upgrading the partner nodes in a batch of HA pairs

You upgrade the partner nodes in a batch of HA pairs by initiating a takeover on the nodes. The first set of nodes serve the nodes' data while the partners are upgraded.

Steps

1. Determine if the nodes to be upgraded are currently serving any clients by entering the following command twice for each node:

   `system node run -node` ***`nodename`*** **`-command uptime`**

   The `uptime` command displays the total number of operations the node has performed for NFS, CIFS, FC, and iSCSI clients since the node was last booted. For each protocol, determine if the operation counts are increasing. If they are increasing, the node is currently serving clients for that protocol. If they are not increasing, the node is not currently serving clients for that protocol.

   You should make a note of each protocol that has increasing client operations so that after the node is upgraded, you can verify that client traffic has resumed.

Example

This example shows a node with NFS, CIFS, FC, and iSCSI operations. However, the node is currently serving only NFS and iSCSI clients.

```
cluster1::> system node run -node node1 -command uptime
  2:58pm up  7 days, 19:16 800000260 NFS ops, 1017333 CIFS ops, 0 HTTP ops, 40395 FCP ops, 32810 iSCSI
ops

cluster1::> system node run -node node1 -command uptime
  2:58pm up  7 days, 19:17 800001573 NFS ops, 1017333 CIFS ops, 0 HTTP ops, 40395 FCP ops, 32815 iSCSI
ops
```

2. Migrate LIFs away from the nodes that will be taken over by entering the following command for each node in the second set:

   `network interface migrate {-data-protocol nfs,cifs -role data -curr-node` ***`source_node`*** **`} -dest-node`** ***`partner_node`***

   Data LIFs for SAN protocols are not migrated. As long as these LIFs exist on each node in the cluster, data can be served through alternate paths during the upgrade process.

3. Verify that the LIFs migrated to the proper ports on the nodes' partners by entering the following command for each node's partner:

   `network interface show -data-protocol nfs,cifs -role data -curr-node` ***`partner_node`***

Example

This example shows that node1's data LIFs were migrated to port e0b on its partner (node0).

```
cluster1::> network interface show -data-protocol nfs,cifs -role data -curr-node node0
            Logical    Status     Network            Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
----------- ---------- ---------- ------------------ ------------- ------- ----
vs0
            lif1       up/up      192.0.2.130/24     node0         e0b     false
            lif2       up/up      192.0.2.131/24     node0         e0b     true
            lif3       up/up      192.0.2.132/24     node0         e0b     false
vs1
            lif1       up/up      192.0.2.133/24     node0         e0b     true
            lif2       up/up      192.0.2.134/24     node0         e0b     false
5 entries were displayed.
```

If desired, you can migrate a LIF to a different port on the partner node by using the `network interface migrate` command.

4. Trigger an AutoSupport notification by entering the following command for each of the nodes in the second set:

   `system node autosupport invoke -node` ***`nodename`*** **`-type all -message "starting_NDU"`**

   This AutoSupport notification includes a record of the system status just prior to upgrade. It saves useful troubleshooting information in case there is a problem with the upgrade process.

   If your cluster is not configured to send AutoSupport messages, a copy of the notification is saved locally.

5. Initiate a takeover by entering one of the following commands for each node in the second set:

If you are upgrading from a...	Enter this command...
Data ONTAP 8.2.x release	**`storage failover takeover -ofnode`** ***`nodename`***
Data ONTAP 8.1.x release	**`storage failover takeover -ofnode`** ***`nodename`*** **`-option allow-version-mismatch`** The `allow-version-mismatch` option enables the HA pair to tolerate different Data ONTAP release family versions during a major release upgrade.

   Do not specify the parameter `-option immediate`, because a normal takeover is required for the node that is being taken over to boot onto the new software image.

   The nodes that are taken over boot up to the `Waiting for giveback` state.

   > **Note:** If AutoSupport is enabled, an AutoSupport message is sent indicating that the nodes are out of cluster quorum. You can safely ignore this notification and proceed with the upgrade.

6. Verify that the takeover was successful:

   `storage failover show`

   Example

   This example shows that the takeover was successful. The second set of nodes (node1, node3, and node5) are in the `Waiting for giveback` state, and their partners are `In takeover`.

```
cluster1::> storage failover show
                              Takeover
Node           Partner        Possible State Description
-------------- -------------- -------- -------------------------------------
node0          node1          false    In takeover
node1          node0          -        Waiting for giveback (HA mailboxes)
node2          node3          false    In takeover
node3          node2          -        Waiting for giveback (HA mailboxes)
node4          node5          false    In takeover
node5          node4          -        Waiting for giveback (HA mailboxes)
node6          node7          true     Connected to node7
node7          node6          true     Connected to node6
node8          node9          true     Connected to node9
node9          node8          true     Connected to node8
node10         node11         true     Connected to node11
node11         node10         true     Connected to node10
12 entries were displayed.
```

7. Wait 8 minutes to ensure the following conditions:
   - Client multipathing (if deployed) is stabilized.
   - Clients are recovered from the pause in I/O that occurs during takeover.

     The recovery time is client-specific and may take longer than 8 minutes depending on the characteristics of the client applications.

8. Return the aggregates to the nodes by entering the following command for each of the nodes' partners:

```
storage failover giveback –ofnode nodename
```

Attention: The giveback is not initiated, an error message is returned, and an event is generated if any conditions such as the following are detected:

- Long-running operations (such as ASUP generation)
- Operations that cannot be restarted (such as aggregate creation)
- Error conditions (such as a disk connectivity mismatch between the nodes)

If giveback is not initiated, complete the following steps:

a. Address the "veto" condition described in the error message, ensuring that any identified operations are terminated gracefully.

b. Reenter the giveback command:

```
storage failover giveback -ofnode nodename
```

Alternatively, you can analyze the messages and events for relevance in your environment. If you determine that the veto conditions are not significant, you can override the giveback veto by entering the following command:

```
storage failover giveback –ofnode nodename -override-vetoes true
```

For more information about determining whether you can safely override the veto, see the *Clustered Data ONTAP High-Availability Configuration Guide*.

This first returns the root aggregate to the partner nodes and then, after those nodes have finished booting, returns the non-root aggregates.

9. Verify that all aggregates have been returned:

```
storage failover show-giveback
```

If the `Giveback Status` field indicates that a node that was taken over is in partial giveback, then complete the following steps before proceeding:

a) Determine which aggregates have been returned:

```
storage aggregate show -node nodename
```

b) Check EMS logs for errors and take corrective action.

For more information about EMS messages, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

c) Repeat Step 8.a to verify any corrections.

The newly-booted nodes begin to serve data to clients from each aggregate as soon as the aggregate is returned.

10. Verify that the upgrade completed successfully for the nodes in the second set:

a) Set the privilege level to advanced:

```
set -privilege advanced
```

b) Ensure that upgrade status is complete for each node in the second set:

`system node upgrade-revert show -node` ***`nodename`***

The status should be listed as `complete`.

If the status is not successful for any node in the set, from the node, run the `system node upgrade-revert upgrade` command. If this command does not complete the node's upgrade, contact technical support immediately.

c) Return to the admin privilege level:

`set -privilege admin`

11. Revert the LIFs back to the second set of nodes:

`network interface revert *`

Example

This command returns the LIFs that were migrated away from the nodes.

```
cluster1::> network interface revert *
12 entries were acted on.
```

12. Verify that the nodes' data ports and LIFs are up and operational by completing the following steps for each node that was upgraded:

a) Verify that the data ports for are up for each node:

`network port show -node` ***`nodename`*** **`-role data`**

Example

This example shows that all of a node's data ports are up.

```
cluster1::> network port show -node node1 -role data
                                      Auto-Negot  Duplex     Speed (Mbps)
Node   Port   Role         Link   MTU Admin/Oper  Admin/Oper Admin/Oper
------ ------ ------------ ---- ----- ----------- ---------- ------------
node1
       e0c    data         up    9000  true/true  full/full   auto/1000
       e0d    data         up    9000  true/true  full/full   auto/1000
       e1b    data         up    9000  true/true  full/full   auto/1000
       e1c    data         up    9000  true/true  full/half   auto/10
       e1d    data         up    9000  true/true  full/half   auto/10
5 entries were displayed.
```

b) Verify that data LIFs successfully reverted back to each node in the first set, and that the LIFs are up:

`network interface show -data-protocol nfs,cifs -role data -curr-node` ***`nodename`***

Example

This example shows that all of the data LIFs hosted by a node have successfully reverted back to the node, and that they are operationally up.

```
cluster1::> network interface show -data-protocol nfs,cifs -role data -curr-node node1
            Logical    Status     Network            Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
----------- ---------- ---------- ------------------ ------------- ------- ----
vs0
            data001      up/up    192.0.2.120/24     node0         e0c     true
            data002      up/up    192.0.2.121/24     node0         e0d     true
            data003      up/up    192.0.2.122/24     node0         e0d     true
            data004      up/up    192.0.2.123/24     node0         e0c     true
4 entries were displayed.
```

13. If you previously determined that any of the partner nodes serve clients, verify that each node is providing service for each protocol that it was previously serving:

`system node run -node` ***`nodename`*** **`-command uptime`**

The operation counts reset to zero during the upgrade.

Example

This example shows that an upgraded node has resumed serving its NFS and iSCSI clients:

```
cluster1::> system node run -node node1 -command uptime
  3:15pm up  0 days, 0:16 129 NFS ops, 0 CIFS ops, 0 HTTP ops, 0 FCP ops, 2 iSCSI ops
```

14. Trigger an AutoSupport notification by entering the following command for each node in the second set:

`system node autosupport invoke -node` ***`nodename`*** **`-type all -message "finishing_NDU"`**

Verifying that the batch was upgraded successfully

After upgrading all of the nodes in a batch, you must verify that the target release is running on the nodes.

Steps

1. Confirm that the new Data ONTAP 8.2.x software is running on all of the nodes in the first batch:

`system node image show`

Example

This example of a 12-node cluster shows version 8.2.1 as the current version on the first batch of nodes (node0 - node5):

```
cluster1::> system node image show
                 Is      Is                 Install
Node     Image   Default Current Version    Date
-------- ------- ------- ------- --------   -------------------
node0
         image1  true    true    8.2.1      11/22/2013 13:52:22
         image2  false   false   8.1.2      10/25/2012 12:37:36
node1
         image1  true    true    8.2.1      11/22/2013 13:52:22
         image2  false   false   8.1.2      10/25/2012 12:41:16
```

```
node2
         image1  true    true    8.2.1      11/22/2013 13:52:22
         image2  false   false   8.1.2      10/25/2012 12:37:36
node3
         image1  true    true    8.2.1      11/22/2013 13:52:22
         image2  false   false   8.1.2      10/25/2012 12:41:16
node4
         image1  true    true    8.2.1      11/22/2013 13:52:22
         image2  false   false   8.1.2      10/25/2012 12:37:36
node5
         image1  true    true    8.2.1      11/22/2013 13:52:22
         image2  false   false   8.1.2      10/25/2012 12:41:16
node6
         image1  true    false   8.2.1      11/22/2013 13:52:22
         image2  false   true    8.1.2      10/25/2012 12:37:36
node7
         image1  true    false   8.2.1      11/22/2013 13:52:22
         image2  false   true    8.1.2      10/25/2012 12:41:16
node8
         image1  true    false   8.2.1      11/22/2013 13:52:22
         image2  false   true    8.1.2      10/25/2012 12:37:36
node9
         image1  true    false   8.2.1      11/22/2013 13:52:22
         image2  false   true    8.1.2      10/25/2012 12:41:16
node10
         image1  true    false   8.2.1      11/22/2013 13:52:22
         image2  false   true    8.1.2      10/25/2012 12:37:36
node11
         image1  true    false   8.2.1      11/22/2013 13:52:22
         image2  false   true    8.1.2      10/25/2012 12:41:16
12 entries were displayed.
```

2. Re-enable automatic giveback on the nodes if it was previously disabled by entering the following command on each node in the first batch:

   `storage failover modify -node` ***`nodename`*** **`-auto-giveback true`**

3. Ensure that the cluster is in quorum and that services are running before upgrading the next batch of nodes.

   You can use the `cluster show` and `cluster ring show` commands to verify that the cluster is in quorum.

After you finish

Upgrade the second batch by following the same steps that you used to upgrade the first batch. After both batches are upgraded, you should verify that the cluster was upgraded successfully.

Upgrading a Data ONTAP cluster disruptively

If you can take your cluster offline to upgrade to a new Data ONTAP release, or if you have a single-node cluster, you can use the disruptive upgrade method. This method has several steps: disabling storage failover for each HA pair, updating the software on each node in the cluster, and then reenabling storage failover.

Before you begin

You must have satisfied upgrade preparation requirements and verified that the cluster is ready to be upgraded.

About this task

During a disruptive upgrade, each node acts as a single-node cluster. Any failures in the node will cause a data outage.

Steps

1. Perform one of the following actions:

If the cluster consists of...	Do this...
One node	Continue to the next step.

Two nodes	**a.** Disable cluster high availability:
	`cluster ha modify -configured false`
	b. Disable storage failover for the HA pair:
	`storage failover modify -node * -enabled false`

More than two nodes	Disable storage failover for each HA pair in the cluster: **`storage failover modify -node * -enabled false`**

2. Reboot each node in the cluster:

   `system node reboot -node` ***`nodename`***

   If the cluster consists of more than one node, you can reboot the nodes simultaneously.

   Each node boots the new Data ONTAP image. The Data ONTAP login prompt appears, indicating that the reboot process is complete.

3. When each node has rebooted with the new Data ONTAP image, confirm that the new Data ONTAP 8.2.x software is running:

   `system node image show`

   #### Example

   This example shows version 8.2.1 as the current version on both nodes:

```
cluster1::> system node image show
                 Is      Is                 Install
Node     Image   Default Current Version    Date
-------- ------- ------- ------- --------   -------------------
node0
         image1  true    true    8.2.1      11/22/2013 13:52:22
         image2  false   false   8.1.2      10/25/2012 12:37:36
node1
         image1  true    true    8.2.1      11/22/2013 13:55:22
         image2  false   false   8.1.2      10/25/2012 12:41:16
4 entries were displayed.
```

4. Verify that the upgrade completed successfully for each node:

   a) Set the privilege level to advanced:

`set -privilege advanced`

b) Ensure that upgrade status is complete for each node:

`system node upgrade-revert show -node` ***`nodename`***

The status should be listed as `complete`.

If the status is not successful, from the node, run the `system node upgrade-revert upgrade` command. If this command does not complete the node's upgrade, contact technical support immediately.

c) Return to the admin privilege level:

`set -privilege admin`

5. Enable storage failover for each HA pair in the cluster:

`storage failover modify -node * -enabled true`

6. If the cluster consists of two nodes, enable cluster high availability:

`cluster ha modify -configured true`

Completing post-upgrade tasks for cluster upgrades

After you upgrade a cluster to the latest version of Data ONTAP software, you must complete additional post-upgrade tasks.

Verifying the cluster version

After all of the HA pairs have been upgraded, you must use the `version` command to verify that all of the nodes are running the target release.

About this task

The cluster version is the lowest version of Data ONTAP running on any node in the cluster.

Step

1. Verify that the cluster version is the target Data ONTAP release:

`version`

Example

```
cluster1::> version
Data ONTAP Release 8.2.1 Cluster-Mode: Fri Nov 22 23:37:32 PDT 2013
(IBM)
```

If the cluster version is not the target Data ONTAP release, run the `system node upgrade-revert upgrade` command to update the cluster version.

Verifying that the cluster is in quorum

Before and after you perform an upgrade, reversion, or downgrade, you must ensure that all nodes are participating in a replicated database (RDB) quorum and that all rings are in the quorum. You must also verify that the per-ring quorum master is the same for all nodes.

About this task

For more information about cluster replication rings and RDB quorums, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

Steps

1. Set the privilege level to advanced:
   `set -privilege advanced`

   Enter **`y`** to continue.
2. Display each RDB process:

To display this RDB process...	Enter this command...
Management application	**`cluster ring show -unitname mgmt`**
Volume location database	**`cluster ring show -unitname vldb`**
Virtual-Interface manager	**`cluster ring show -unitname vifmgr`**
SAN management daemon	**`cluster ring show -unitname bcomd`**

Example

This example shows the volume location database process for a cluster running Data ONTAP 8.1.x:

```
cluster1::*> cluster ring show -unitname vldb

Node   UnitName Epoch DB Epoch DB Trnxs Master
------ -------- ----- -------- -------- ---------
node0  vldb     154   154      14847    node0
node1  vldb     154   154      14847    node0
node2  vldb     154   154      14847    node0
node3  vldb     154   154      14847    node0
4 entries were displayed.
```

Example

This example shows the volume location database process for a cluster running Data ONTAP 8.2.x:

```
cluster1::*> cluster ring show -unitname vldb
Node      UnitName Epoch    DB Epoch DB Trnxs Master    Online
```

```
--------- -------- -------- -------- -------- --------- ---------
node0     vldb     154      154      14847    node0     master
node1     vldb     154      154      14847    node0     secondary
node2     vldb     154      154      14847    node0     secondary
node3     vldb     154      154      14847    node0     secondary
4 entries were displayed.
```

For each process, verify the following configuration details:

- The relational database epoch and database epochs match for each node.
- The per-ring quorum master is the same for all nodes.
  Note that each ring might have a different quorum master.

3. Return to the admin privilege level:
   `set -privilege admin`
4. If you are operating in a SAN environment, verify that each node is in a SAN quorum:
   `event log show -messagename scsiblade.*`

   The most recent `scsiblade` event message for each node should indicate that the scsi-blade is in quorum. During the upgrade, reversion, or downgrade process, each node will temporarily fall out of SAN quorum. Therefore, if you are verifying the SAN quorum after completing an upgrade, reversion, or downgrade, you may notice critical event messages warning you that the nodes were previously out of SAN quorum.

   If a node is out of SAN quorum, you can use the `storage failover takeover` and `storage failover giveback` commands to perform a planned takeover and giveback with the node's high-availability partner and bring the node back into SAN quorum.

Example

In this example, both nodes in the cluster are in SAN quorum.

```
cluster1::> event log show -messagename scsiblade.*
Time                Node             Severity      Event
------------------- ---------------- ------------- ---------------------------
8/13/2013 14:03:51  node0            INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
8/13/2013 14:03:51  node1            INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
```

Example

This example shows a two-node cluster after performing an upgrade. Each node shows a previous out of SAN quorum event message from when the node was upgraded. However, the most recent event message for each node shows that both nodes are presently in SAN quorum.

```
cluster1::> event log show -messagename scsiblade.*
Time                Node             Severity      Event
------------------- ---------------- ------------- ---------------------------
8/13/2013 15:37:51  node1            INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
8/13/2013 15:31:26  node1            CRITICAL      scsiblade.out.of.quorum: This node ...
8/13/2013 15:31:26  node1            INFORMATIONAL scsiblade.ics.trace.dumpfile: The ...
8/13/2013 15:31:26  node1            INFORMATIONAL scsiblade.ics.trace.dumpfile: The ...
8/13/2013 15:30:43  node0            INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
8/13/2013 15:24:16  node0            CRITICAL      scsiblade.out.of.quorum: This node ...
8/13/2013 15:24:16  node0            INFORMATIONAL scsiblade.ics.trace.dumpfile: The ...
8/13/2013 15:24:16  node0            INFORMATIONAL scsiblade.ics.trace.dumpfile: The ...
```

Verifying cluster and SVM health

Before and after you upgrade, revert, or downgrade a cluster, you should verify that the nodes are healthy and eligible to participate in the cluster, and the aggregates and volumes are online.

Steps

1. Verify that the nodes in the cluster are online and are eligible to participate in the cluster:

   `cluster show`

   #### Example

   ```
   cluster1::> cluster show
   Node                  Health  Eligibility
   --------------------- ------- ------------
   node0                 true    true
   node1                 true    true
   ```

   If any node is unhealthy or ineligible, check EMS logs for errors and take corrective action. For more information about EMS messages, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

2. Determine if any disk drives are broken, undergoing maintenance, or reconstructing:

To check for...	Do this...
Broken disks	**a.** Display any broken disks: **`storage disk show -state broken`** **b.** Remove or replace any broken disks.
Disks undergoing maintenance or reconstructing	**a.** Display any disks in maintenance, pending, or reconstructing states: **`storage disk show -state maintenance|pending|reconstructing`** **b.** Wait for the maintenance or reconstruction operation to complete before proceeding.

3. To verify that all aggregates are online, display the state of physical and logical storage, including storage aggregates:

   `storage aggregate show -state !online`

   This command displays the aggregates that are *not* online.

   #### Example

   ```
   cluster1::> storage aggregate show -state !online
   There are no entries matching your query.
   ```

For more information about managing aggregates, see the *Clustered Data ONTAP Physical Storage Management Guide*.

4. To verify that all volumes are online, display any volumes *not* online:

   `volume show -state !online`

   Example

   ```
   cluster1::> volume show -state !online
   There are no entries matching your query.
   ```

   For more information about managing volumes, see the *Clustered Data ONTAP Logical Storage Management Guide*.

Enabling and reverting LIFs to home ports

During a reboot, some LIFs might have been migrated to their assigned failover ports. Before and after you upgrade, revert, or downgrade a cluster, you must enable and revert any LIFs that are not on their home ports.

About this task

The `network interface revert` command reverts a LIF that is not currently on its home port back to its home port, provided that the home port is operational. A LIF's home port is specified when the LIF is created; you can determine the home port for a LIF by using the `network interface show` command.

Steps

1. Display the status of all LIFs:

   `network interface show`

   Example

   This example displays the status of all LIFs for a Storage Virtual Machine (SVM, formerly known as Vserver).

   ```
   cluster1::> network interface show -vserver vs0
               Logical    Status     Network            Current       Current Is
   Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
   ----------- ---------- ---------- ------------------ ------------- ------- ----
   vs0
               data001    down/down  192.0.2.120/24     node0         e0e     true
               data002    down/down  192.0.2.121/24     node0         e0f     true
               data003    down/down  192.0.2.122/24     node0         e2a     true
               data004    down/down  192.0.2.123/24     node0         e2b     true
               data005    down/down  192.0.2.124/24     node0         e0e     false
               data006    down/down  192.0.2.125/24     node0         e0f     false
               data007    down/down  192.0.2.126/24     node0         e2a     false
               data008    down/down  192.0.2.127/24     node0         e2b     false
   8 entries were displayed.
   ```

If any LIFs appear with a `Status Admin` status of `down` or with an `Is home` status of `false`, continue with the next step.

2. Enable the data LIFs:

   `network interface modify {-role data} -status-admin up`

   Example

   ```
   cluster1::> network interface modify {-role data} -status-admin up
   8 entries were modified.
   ```

3. Revert LIFs to their home ports:

   `network interface revert *`

   Example

   This command reverts all LIFs back to their home ports and changes all LIF home statuses to `true`.

   ```
   cluster1::> network interface revert *
   8 entries were acted on.
   ```

4. Verify that all LIFs are in their home ports:

   `network interface show`

   Example

   This example shows that all LIFs for SVM vs0 are on their home ports.

   ```
   cluster1::> network interface show -vserver vs0
               Logical    Status     Network            Current       Current Is
   Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
   ----------- ---------- ---------- ------------------ ------------- ------- ----
   vs0
               data001      up/up    192.0.2.120/24     node0         e0e     true
               data002      up/up    192.0.2.121/24     node0         e0f     true
               data003      up/up    192.0.2.122/24     node0         e2a     true
               data004      up/up    192.0.2.123/24     node0         e2b     true
               data005      up/up    192.0.2.124/24     node1         e0e     true
               data006      up/up    192.0.2.125/24     node1         e0f     true
               data007      up/up    192.0.2.126/24     node1         e2a     true
               data008      up/up    192.0.2.127/24     node1         e2b     true
   8 entries were displayed.
   ```

Creating a namespace mirror constituent for an upgraded Infinite Volume

After upgrading a cluster that contains an Infinite Volume, you must ensure that a namespace mirror constituent is created on all Infinite Volumes that span two or more nodes to provide data protection for the namespace constituent.

Before you begin

- An upgraded Infinite Volume must span two or more nodes to create a namespace mirror constituent.
- Determine whether to use this procedure or contact technical support.
  - If technical support created a data protection mirror copy for the namespace constituent in an Infinite Volume running Data ONTAP 8.1.x before you upgraded, you must contact technical support after the upgrade to convert the data protection mirror copy of the namespace constituent to a namespace mirror constituent.
  - If technical support did not create a data protection mirror copy for the namespace constituent in an Infinite Volume running Data ONTAP 8.1.x before you upgraded, you can use this procedure to create a namespace mirror constituent after the upgrade.

About this task

All the Infinite Volumes in a cluster running Data ONTAP 8.2 or later must have a namespace mirror constituent. An Infinite Volume that was upgraded from Data ONTAP 8.1.x does not have a namespace mirror constituent.

Note: When an Infinite Volume is in a data protection mirror relationship, only the source read/write Infinite Volume requires a namespace mirror constituent to provide data protection for the namespace constituent. The destination read-only Infinite Volume does not require a namespace mirror constituent because the data protection mirror relationship for the Infinite Volume provides data protection.

Steps

1. Ensure that the upgraded Infinite Volume has enough space for a namespace mirror constituent.

   a) View the size of the namespace constituent on the Infinite Volume, and view the name of the aggregate that contains the namespace constituent by using the `volume show` command with the `-is-constituent true` parameter.

      You need the size of the namespace constituent because the namespace mirror constituent will be the same size as the namespace constituent. You need the name of the aggregate because you must eliminate this aggregate from your evaluation. The namespace mirror constituent must be on a different aggregate than the namespace constituent, and the namespace mirror constituent should be on a different node than the namespace constituent.

Example

In the following example, the namespace constituent is named repo_vol_ns, and it is 10 TB in size. The name of the aggregate that contains the namespace constituent is aggr_ns.

```
cluster1::> volume show -is-constituent true
Vserver Volume                 Aggregate State  Type Size  Available Used%
------- ------------           --------- ------ ---- ----- --------- -----
vs0     repo_vol_default_data0001 aggr1   online RW   50TB    35.0TB   30%
vs0     repo_vol_default_data0002 aggr2   online RW   50TB    34.0TB   32%
vs0     repo_vol_default_data0003 aggr3   online RW   50TB    35.5TB   29%
vs0     repo_vol_default_data0004 aggr4   online RW   50TB    36.5TB   27%
vs0     repo_vol_ns               aggr_ns online RW   10TB     8.4TB   16%
5 entries were displayed.
```

b) In the cluster with the Infinite Volume, find an aggregate that is not on the same node as the aggregate that contains the namespace constituent by using the `aggregate show` command with the `-fields node` parameters.

Example

In the following example, the aggregate that contains the namespace constituent is named aggr_ns, and the aggregate is on node-02. You should not use any of the following aggregates on node-02: aggr1 or aggr_ns. Instead you should use one of the following aggregates on node-01: aggr2, aggr3, or aggr4.

```
cluster1::> aggregate show -fields node
aggregate node
--------- ------------
aggr1     node-02
aggr2     node-01
aggr3     node-01
aggr4     node-01
aggr_ns   node-02
```

c) Ensure one of the aggregates has enough space to fit the namespace mirror constituent.

d) Increase aggregate space if required.

The Infinite Volume has enough space for the namespace mirror constituent.

2. Increase the size of the Infinite Volume by the size of the namespace constituent by using the `volume modify` command.

Example

In the following example, the volume named repo_vol is increased by the size of the namespace constituent, which is 10 TB:

```
cluster1::> volume modify -vserver vs0 -volume repo_vol -size +10TB
```

A namespace mirror constituent is automatically created.

Setting the cluster management LIF for Remote Support Agent

After you perform a Data ONTAP upgrade, if you use Remote Support Agent (RSA), you must use the `rsa setup` command on each remote management device to set the cluster management LIF for the RSA software.

Resuming SnapMirror operations

After completing a nondisruptive upgrade or downgrade, you must resume any SnapMirror relationships that were suspended.

Before you begin

Existing SnapMirror relationships must have been suspended by using the `snapmirror quiesce` command, and the cluster must have been nondisruptively upgraded or downgraded.

Steps

1. Resume transfers for each SnapMirror relationship that was previously quiesced:

   `snapmirror resume *`

2. Verify that the SnapMirror operations have resumed:

   `snapmirror show`

Example

```
cluster1::> snapmirror show

Source            Destination  Mirror  Relationship  Total            Last
Path        Type  Path         State   Status        Progress  Healthy Updated
----------- ---- ------------ ------- -------------- --------- ------- --------
cluster1-vs1:dp_src1
            DP   cluster1-vs2:dp_dst1
                              Snapmirrored
                                      Idle           -         true    -
cluster1-vs1:xdp_src1
            XDP  cluster1-vs2:xdp_dst1
                              Snapmirrored
                                      Idle           -         true    -
cluster1://cluster1-vs1/ls_src1
            LS   cluster1://cluster1-vs1/ls_mr1
                              Snapmirrored
                                      Idle           -         true    -
                 cluster1://cluster1-vs1/ls_mr2
                              Snapmirrored
                                      Idle           -         true    -
4 entries were displayed.
```

For each SnapMirror relationship, verify that the Relationship Status is "Idle". If the status is "Transferring", wait for the SnapMirror transfer to complete, and then reenter the command to verify that the status has changed to "Idle".

After you finish

For each SnapMirror relationship that is configured to run on a schedule, you should verify that the first scheduled SnapMirror transfer completes successfully.

Upgrading older style data protection mirror relationships

Data protection mirror relationships that you have before upgrading to Data ONTAP 8.2 use a different relationship syntax and are not upgraded when you upgrade Data ONTAP. You must upgrade relationship syntax style to Data ONTAP 8.2 syntax if you want to use new features and improvements.

Steps

1. Show the relationship-capability of the data protection mirror relationships by using the `snapmirror show` command with the `-fields relationship-capability` parameter.
2. Create SVM peer relationships:
   a) On one of the clusters of the cluster peer relationship, create SVM peer relationships by using the `vserver peer create` command.

      If you have more than one SnapMirror relationship that uses the same two SVMs, then you need to create only one SVM peer relationship.

   b) On the other cluster of the cluster peer relationship, accept the SVM create request by using the `vserver peer accept` command.

   c) If there is a name conflict, change the name of one of the SVMs by using the `vserver rename` command.

   In addition to creating a SVM peer relationship, Data ONTAP upgrades the data protection mirror relationships to Data ONTAP 8.2 style relationships.

   Note: Load-sharing mirror relationships cannot be upgraded to the Data ONTAP 8.2 style. The “Relationship capability” field of the detailed view of the `snapmirror show` command will show as “Pre 8.2”.

3. Ensure that all of the data protection mirror relationships are displayed as `healthy` and have relationship-capability of `8.2 and above` by using the `snapmirror show` command with the `-instances` parameter.

Enabling automatic LIF rebalancing

If you previously disabled automatic LIF rebalancing to perform a batch upgrade, you should reenable it after completing the upgrade.

Steps

1. Set the privilege level to advanced:

   `set -privilege advanced`

2. Enable automatic LIF rebalancing for each LIF as needed:

   ```
   network interface modify -vserver Vserver_name -lif LIF_name -allow-lb-migrate true
   ```

3. Return to the admin privilege level:

   ```
   set -privilege admin
   ```

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 cluster. You might also need to update firmware between Data ONTAP upgrades.

How system firmware is updated during the Data ONTAP upgrade

When you perform a Data ONTAP software upgrade, the firmware service image included with the Data ONTAP upgrade package is copied to each node's boot device, and the new firmware is installed automatically.

If you are upgrading system firmware between Data ONTAP upgrades, you can obtain system firmware and information about how to install it from the N series support website (accessed and navigated as described in *Websites* on page 7).

Disk firmware updates

Disk firmware is bundled with the Data ONTAP upgrade package and updated automatically during Data ONTAP upgrades. You can also obtain disk firmware from the N series support website (accessed and navigated as described in *Websites* on page 7) and update it manually.

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 upgrade package. You can also update disk firmware by downloading the most recent firmware package from the N series support website (accessed and navigated as described in *Websites* on page 7) and installing the files.

Each node is shipped with 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 any firmware bundle resident on the node.

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

- Data ONTAP detects disk firmware updates on the node.
  Data ONTAP scans for new firmware every two minutes.

Disk firmware updates can be added to the node 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 website (accessed and navigated as described in *Websites* on page 7) if 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.
- When you hot-add a SAS shelf

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 cluster 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.

When you need to update the Disk Qualification Package

The Disk Qualification Package (DQP) adds full support for newly qualified disk drives. Each time you update disk firmware or add new disk types or sizes to the cluster, you also need to update the DQP.

You can obtain the DQP from the N series support website (accessed and navigated as described in *Websites* on page 7). You need to download and install the DQP in the following situations:

- Whenever you add a new disk type or size to the node

  For example, if you already have 1-TB disks and add 2-TB disks, you need to check for the latest DQP update.
- Whenever you update the disk firmware
- Whenever newer disk firmware or DQP files are available

Related information

Disk Qualification Package Instructions: www.ibm.com/support/docview.wss?uid=ssg1S7002516

Disk Drive & Firmware Matrix: www.ibm.com/support/docview.wss?uid=ssg1S7002517

Service availability during disk firmware updates

By default, disk firmware updates take place automatically in the background so as to ensure the continuity of cluster services.

You can download the disk firmware package to your cluster at any time and the firmware is updated nondisruptively in the background. However, you must wait until the disk firmware update has finished before initiating a nondisruptive upgrade.

Background disk firmware updates take place one disk at a time and require approximately 2.5 minutes per disk. Although it is not likely that all the disks attached to your cluster would need firmware updates at the same time, it is a best practice to wait at least 2.5 minutes for every disk attached to a node before proceeding with a Data ONTAP NDU.

For example, if a node has 192 disks attached, you should wait at least 480 minutes, or 8 hours. You should wait until all nodes in the cluster have completed the firmware update before proceeding with the NDU. Disk firmware can be updated in parallel on all nodes in a cluster.

Disk shelf firmware updates

Disk shelf firmware (firmware for modules on disk shelves) is bundled with the Data ONTAP upgrade package and updated automatically during Data ONTAP upgrades. You can also obtain disk shelf firmware from the N series support website (accessed and navigated as described in *Websites* on page 7) and update it manually.

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

How disk shelf firmware is updated

When you upgrade Data ONTAP, disk shelf firmware is updated automatically if the firmware on the shelves is older than the firmware that is bundled with the Data ONTAP upgrade package. You can also update disk shelf firmware by downloading and installing the most recent firmware for your shelf modules from the N series support website (accessed and navigated as described in *Websites* on page 7).

The AT series, ESH series, and SAS shelf I/O module (IOM) series in a disk shelf provide for the connection 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 certain systems. Updated firmware for these modules is made available periodically.

Each node is shipped with the latest disk shelf firmware versions.

Disk shelf firmware updates can be loaded onto the cluster at the following times:

- After a Data ONTAP upgrade
  Disk shelf firmware updates are often included in Data ONTAP upgrade packages. After the Data ONTAP upgrade process is completed on each node in an HA pair, if the new disk shelf firmware version is later than the installed version, the new version is downloaded and installed on the disk shelves attached to both nodes in the HA pair.
- During a manual firmware update
  You might need to download a disk shelf firmware update from the N series support website (accessed and navigated as described in *Websites* on page 7) 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

The following events can also trigger an automatic disk shelf firmware update when there is new firmware on the cluster:

- The `system reboot` command is issued.
- The `storage failover giveback` command is issued.

- New disk drives are inserted.
- New shelf modules are inserted.

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

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 determine which firmware version is installed on disk shelves attached to your cluster.

Steps

1. Go to the disk shelf firmware information on the N series support website (accessed and navigated as described in *Websites* on page 7) and determine the most recent firmware version for your shelves.
2. At the clustershell, enter the following command:

   `system node run` ***`nodename`*** **`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 website (accessed and navigated as described in *Websites* on page 7), you must update your disk shelf firmware manually.

How ACP firmware is updated

If your disk shelves include ACP functionality, ACP firmware is updated automatically during Data ONTAP upgrades. You can also obtain it from the N series support website (accessed and navigated as described in *Websites* on page 7) and update it manually.

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 upgrade package.

Service Processor firmware updates

Service Processor (SP) is a remote management device that is included in some clusters. Starting with Data ONTAP 8.2, a baseline SP firmware image is packaged with the Data ONTAP upgrade package, so that the SP is updated automatically by default.

Any existing connection to the SP is terminated when the SP firmware is being updated. This is the case whether the SP firmware update is automatic or manually triggered.

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

If you are upgrading SP firmware between Data ONTAP upgrades, you can obtain SP firmware and information about how to install it from the N series support website (accessed and navigated as described in *Websites* on page 7). You can download and update the SP firmware by using the Data ONTAP CLI or the SP CLI.

For information about what the SP is and how it works, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

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.

You can obtain RLM firmware and information about how to install it from the N series support website (accessed and navigated as described in *Websites* on page 7).

For information about what the RLM is and how it works, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

How Flash Cache firmware is updated

Firmware for Flash Cache devices is included with the upgrade package for Data ONTAP upgrades. If the running firmware is older than the firmware that is bundled with the Data ONTAP upgrade package, it is updated automatically.

Firmware updates are not available for the original 16-GB PAM devices. Automatic updates occur only to Flash Cache devices, not PAM 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 `storage failover giveback` process. Consequently, if your cluster 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 *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

Reverting clusters to an earlier Data ONTAP release family

Transitioning a cluster to a release in an earlier Data ONTAP family is referred to as a *reversion*. Reverting requires preparation, using the `system node revert-to` command in the clustershell, using the `revert_to` command in the nodeshell, 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 cluster 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 clustershell `system node revert-to` and nodeshell `revert_to` commands for the reversion process.

For more information, see the `system node revert-to` man page.

When to revert and when to call technical support

You can revert without assistance when reverting new or test clusters, but you should call technical support if you encounter problems during or after upgrading, or if you want to revert a production cluster.

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 cluster and you want to return to the original release when testing is completed.
- You are configuring a new cluster—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 in a production environment without assistance. If you encounter any of the following circumstances, contact technical support immediately:

- The upgrade process fails and cannot finish.
- The upgrade process finishes, but the cluster is unusable in a production environment.
- The upgrade process finishes and the cluster goes into production, but you are not satisfied with its behavior.
- The upgrade process finishes for some but not all of the nodes, and you decide that you want to revert.

Planning your reversion

Because new features are introduced in each release of Data ONTAP, you must understand reversion requirements and evaluate how they might impact your current configuration.

Before proceeding with the reversion, you should do the following:

- Review the *Release Notes* for the Data ONTAP reversion source release.
- Understand any requirements for reverting to the target release from your existing software.
- Note any potential functionality changes to your cluster after the reversion.
- Prepare to address all points in the reversion checklist.

Cluster reversion checklist

You can use this checklist to record your progress as you prepare for the reversion, perform the reversion, and complete post-reversion tasks.

Steps for preparing to revert

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

Condition	**Complete?**
Software and hardware support in the target release is confirmed. To confirm hardware support, visit *IBM System Storage N series Introduction and Planning Guide* at the N series support website (accessed and navigated as described in *Websites* on page 7) If any nodes are platforms that are not supported in the target release, then you should unjoin the nodes from the cluster before you revert the cluster. For more information about removing nodes from a cluster, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.	
The SAN configuration is fully supported. You must verify that your SAN configuration does not exceed the limits for the target Data ONTAP release. For more information about SAN configuration limits, see the *Clustered Data ONTAP SAN Configuration Guide*.	
All release-specific reversion issues have been resolved.	
You have access to the clustershell at the advanced privilege level.	

Condition	Complete?
The cluster and Storage Virtual Machines (SVMs) are running and healthy. All aggregates and volumes should be healthy and online before proceeding with the reversion. You can use the `cluster show` command to verify the status of the nodes.	
The cluster is in quorum. All nodes are participating in a quorum and all rings are in the quorum. The per-ring quorum master should be the same for all nodes.	
All core dump files have been either saved or deleted. For more information about managing core dumps, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.	
The system time is synchronized across the cluster.	
The iSNS server and its associated SVM management LIF are configured with IPv4 addresses.	
Snapshot copies created in the later Data ONTAP release have been deleted, and all Snapshot copies from the nodes' root volumes and root aggregates have been deleted. You must delete any Snapshot copies that were created after upgrading to the current release, delete root aggregate and root volume Snapshot copies, and disable Snapshot schedules for the root aggregate and root volume.	
All IPv6 objects have been deleted. You must delete all LIFs, routing groups, and firewall policies that use IPv6 addresses. For more information about deleting these objects, see the *Clustered Data ONTAP Network Management Guide*.	
You have the target Data ONTAP software image available on an HTTP server. Download the software image for the target Data ONTAP release from the N series support website (accessed and navigated as described in *Websites* on page 7), and place it on an HTTP server that is accessible by each node.	

Condition	Complete?
The target Data ONTAP software images are installed on each node and set as the alternate boot device image. You can use the `system node image update` command to install the software images. You can use the `system node image show` command to verify that the software images are installed as the alternate boot image on each node.	

Steps for performing the reversion

The reversion is complete when all of the following steps have been completed:

Condition	Complete?
No jobs are running. If any aggregate, volume, mirror, NDMP (dump or restore), or Snapshot jobs (such as create, delete, move, modify, replicate, and mount jobs) are running or queued, allow the jobs to complete successfully or stop the queued entries.	
The target Data ONTAP software is installed and set as the default boot image.	
Each node has booted the target release.	

Steps for after reverting

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

Condition	Complete?
The cluster and SVMs are running and healthy. All aggregates and volumes should be healthy and online after the reversion. You can use the `cluster show` command to verify the status of the nodes.	
LIFs are online and on their correct home ports. You can use the `network interface` command to display and modify LIF configuration.	
Snapshot schedules are enabled for each node's root volume and root aggregate. You must enable Snapshot schedules of the root volume and root aggregate to start creating Snapshot copies again.	

Condition	**Complete?**
Client access is verified. You should verify that clients can access the cluster for each configured protocol.	
The Service Processor firmware version is confirmed. If your current SP firmware version is not supported for the Data ONTAP release to which you reverted, you must install a supported SP firmware version for the earlier Data ONTAP release.	
The event generate-autosupport-log capability is enabled for the SVM administrator vsadmin role.	

Reversion process considerations

Make sure that you know about revert issues and limitations before beginning a Data ONTAP reversion.

Be aware of the following issues:

- Clusters consisting of more than one node can be reverted from Data ONTAP 8.2.x to 8.1.x.

  Single node clusters cannot be reverted to Data ONTAP 8.1.
- If the cluster contains a Storage Virtual Machine (SVM) with Infinite Volume, do not attempt to revert to Data ONTAP 8.1.

  If you are reverting to Data ONTAP 8.1.1 or a later release of Data ONTAP 8.1.x, you should contact technical support for assistance. If you are reverting to Data ONTAP 8.1.0, you must delete all SVMs with Infinite Volume.
- Reversion is disruptive.

  No client access can occur during the reversion. If you are reverting a production cluster, be sure to include this disruption in your planning.
- Reversion affects all nodes in the cluster.

  The reversion must be performed on all nodes in the cluster; however, some of the procedures must be performed on each HA pair and completed on each set of nodes before other pairs are reverted.
- The reversion is complete when all nodes are running the new target release.

  When the cluster is in a mixed-version state, you should not enter any commands that alter the cluster operation or configuration except as necessary to satisfy reversion requirements; monitoring operations are permitted.

  If you are unable to complete the reversion for any reason, contact technical support immediately. If you have reverted some, but not all, of the nodes, do not attempt to upgrade the cluster back to the source release.
- You might notice reduced performance for client operations.

  Storage QoS, which prioritizes client and system operations, is not available when you revert from Data ONTAP 8.2 to 8.1.x.

- When you revert a node, it clears the cached data in a Flash Cache module.

  Because there is no cached data in the Flash Cache module, the node serves initial read requests from disk, which results in decreased read performance during this period. The node repopulates the cache as it serves read requests.
- When you revert a volume containing FlexCache volumes from Data ONTAP 8.2 to Data ONTAP 8.1.x, the FlexCache volumes are not deleted.

  The FlexCache volumes become dormant and cannot be used for caching read requests from clients.
- After you enter the `system node revert-to` command to revert a cluster, the `version` command becomes unavailable and does not display any output until the reversion is completed.

Identifying potential reversion issues

Every Data ONTAP release family has unique reversion requirements that you must understand and resolve before you decide to revert.

For additional information, and to check for reversion issues that might have been discovered later, see the *Clustered Data ONTAP Release Notes*. The following list summarizes reversion issues known when this guide was published.

- If you gained access to licensed or entitled features by upgrading to Data ONTAP 8.2 or later and these features required a license prior to Data ONTAP 8.2, you might need to install a license after reverting.

  This is the case if a license was not previously installed for the release to which you revert.
- You need to be aware of the licensing implications if you have a node that is shipped with Data ONTAP 8.2 and you want to revert it to a release in the Data ONTAP 8.1 release family.

  If your system is shipped with Data ONTAP 8.2 or later and you reinstall it with a release in the Data ONTAP 8.1 release family, you must install a key for the desired license package in the format supported by the Data ONTAP 8.1 release family.
- If your current SP firmware version is not supported for the Data ONTAP release to which you are downgrading or reverting, you must install a supported SP firmware version for the earlier Data ONTAP release.
- Support for up to 1,024 ACEs in NFSv4 ACLs is a new feature in Data ONTAP 8.2.

  Reverting to a previous release family requires action if you have NFSv4 ACLs in your environment that contain more than 400 ACEs.
- Before reverting to a release earlier than Data ONTAP 8.2.1, you must ensure that off-board antivirus is disabled on all the nodes within a cluster.

  If you do not disable off-board antivirus on all the nodes, you will not be able to revert to an earlier release.
- If you downgrade or revert to a release earlier than Data ONTAP 8.2.1, you must run the `security ssh prepare-to-downgrade` command at the advanced privilege level to reset the SSH security configurations of the cluster and all Storage Virtual Machines (SVMs) to the default settings used in the earlier release.
- Exporting qtrees is not supported in releases earlier than Data ONTAP 8.2.1.

If you plan to downgrade or revert to an earlier release, you must first take certain actions; otherwise, the downgrade or revert operation fails.

- Starting with Data ONTAP 8.2, NTP is always enabled.

  If you revert to an earlier release, NTP remains enabled regardless of the prior setting for NTP when you were running the earlier release.
- Before reverting to a release earlier than Data ONTAP 8.2, you must remove the non-server types of digital certificates and the user logins that have the `cert` authentication method.
- If you revert to Data ONTAP 8.1.2 or earlier, Data ONTAP prompts you to run the `security login role config reset` command to reset some role-based access control (RBAC) characteristics to their default values.
- RSH and IPv6 connections are supported only for Data ONTAP 8.2 and later releases.

  Before reverting to an earlier release, you must manually remove RSH user accounts and the IPv6 firewall policies that were added for enabling clear-text protocols (RSH and Telnet).
- AD domain users' access to the cluster is supported only on Data ONTAP 8.1.1 and later releases.

  If you revert to an earlier release, Data ONTAP prompts you to delete any existing authentication tunnel that is used for authenticating AD domain users' cluster access.

Preparing to revert Data ONTAP clusters

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.

Verifying that the cluster is in quorum

Before and after you perform an upgrade, reversion, or downgrade, you must ensure that all nodes are participating in a replicated database (RDB) quorum and that all rings are in the quorum. You must also verify that the per-ring quorum master is the same for all nodes.

About this task

For more information about cluster replication rings and RDB quorums, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

Steps

1. Set the privilege level to advanced:

   `set -privilege advanced`

   Enter **`y`** to continue.
2. Display each RDB process:

To display this RDB process...	Enter this command...
Management application	**`cluster ring show -unitname mgmt`**
Volume location database	**`cluster ring show -unitname vldb`**
Virtual-Interface manager	**`cluster ring show -unitname vifmgr`**
SAN management daemon	**`cluster ring show -unitname bcomd`**

Example

This example shows the volume location database process for a cluster running Data ONTAP 8.1.x:

```
cluster1::*> cluster ring show -unitname vldb

Node   UnitName Epoch DB Epoch DB Trnxs Master
------ -------- ----- -------- -------- ---------
node0  vldb     154   154      14847    node0
node1  vldb     154   154      14847    node0
node2  vldb     154   154      14847    node0
node3  vldb     154   154      14847    node0
4 entries were displayed.
```

Example

This example shows the volume location database process for a cluster running Data ONTAP 8.2.x:

```
cluster1::*> cluster ring show -unitname vldb
Node      UnitName Epoch    DB Epoch DB Trnxs Master    Online
--------- -------- -------- -------- -------- --------- ---------
node0     vldb     154      154      14847    node0     master
node1     vldb     154      154      14847    node0     secondary
node2     vldb     154      154      14847    node0     secondary
node3     vldb     154      154      14847    node0     secondary
4 entries were displayed.
```

For each process, verify the following configuration details:

- The relational database epoch and database epochs match for each node.
- The per-ring quorum master is the same for all nodes.
  Note that each ring might have a different quorum master.

3. Return to the admin privilege level:
   `set -privilege admin`
4. If you are operating in a SAN environment, verify that each node is in a SAN quorum:
   `event log show -messagename scsiblade.*`

   The most recent `scsiblade` event message for each node should indicate that the scsi-blade is in quorum. During the upgrade, reversion, or downgrade process, each node will temporarily fall out of SAN quorum. Therefore, if you are verifying the SAN quorum after completing an upgrade, reversion, or downgrade, you may notice critical event messages warning you that the nodes were previously out of SAN quorum.

If a node is out of SAN quorum, you can use the `storage failover takeover` and `storage failover giveback` commands to perform a planned takeover and giveback with the node's high-availability partner and bring the node back into SAN quorum.

Example

In this example, both nodes in the cluster are in SAN quorum.

```
cluster1::> event log show -messagename scsiblade.*
Time                Node             Severity      Event
------------------- ---------------- ------------- --------------------------
8/13/2013 14:03:51  node0            INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
8/13/2013 14:03:51  node1            INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
```

Example

This example shows a two-node cluster after performing an upgrade. Each node shows a previous out of SAN quorum event message from when the node was upgraded. However, the most recent event message for each node shows that both nodes are presently in SAN quorum.

```
cluster1::> event log show -messagename scsiblade.*
Time                Node             Severity      Event
------------------- ---------------- ------------- --------------------------
8/13/2013 15:37:51  node1            INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
8/13/2013 15:31:26  node1            CRITICAL      scsiblade.out.of.quorum: This node ...
8/13/2013 15:31:26  node1            INFORMATIONAL scsiblade.ics.trace.dumpfile: The ...
8/13/2013 15:31:26  node1            INFORMATIONAL scsiblade.ics.trace.dumpfile: The ...
8/13/2013 15:30:43  node0            INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
8/13/2013 15:24:16  node0            CRITICAL      scsiblade.out.of.quorum: This node ...
8/13/2013 15:24:16  node0            INFORMATIONAL scsiblade.ics.trace.dumpfile: The ...
8/13/2013 15:24:16  node0            INFORMATIONAL scsiblade.ics.trace.dumpfile: The ...
```

Verifying cluster and SVM health

Before and after you upgrade, revert, or downgrade a cluster, you should verify that the nodes are healthy and eligible to participate in the cluster, and the aggregates and volumes are online.

Steps

1. Verify that the nodes in the cluster are online and are eligible to participate in the cluster:

   `cluster show`

Example

```
cluster1::> cluster show
Node                  Health  Eligibility
--------------------- ------- ------------
node0                 true    true
node1                 true    true
```

If any node is unhealthy or ineligible, check EMS logs for errors and take corrective action. For more information about EMS messages, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

2. Determine if any disk drives are broken, undergoing maintenance, or reconstructing:

To check for...	Do this...
Broken disks	**a.** Display any broken disks: **`storage disk show -state broken`** **b.** Remove or replace any broken disks.
Disks undergoing maintenance or reconstructing	**a.** Display any disks in maintenance, pending, or reconstructing states: **`storage disk show -state maintenance|pending|reconstructing`** **b.** Wait for the maintenance or reconstruction operation to complete before proceeding.

3. To verify that all aggregates are online, display the state of physical and logical storage, including storage aggregates:

   `storage aggregate show -state !online`

   This command displays the aggregates that are *not* online.

   Example

   ```
   cluster1::> storage aggregate show -state !online
   There are no entries matching your query.
   ```

   For more information about managing aggregates, see the *Clustered Data ONTAP Physical Storage Management Guide*.

4. To verify that all volumes are online, display any volumes *not* online:

   `volume show -state !online`

   Example

   ```
   cluster1::> volume show -state !online
   There are no entries matching your query.
   ```

   For more information about managing volumes, see the *Clustered Data ONTAP Logical Storage Management Guide*.

Preparing to revert production clusters

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

Ending SnapMirror and SnapVault relationships before reverting

When you revert a cluster, you must end SnapMirror and SnapVault relationships before reverting. Depending on the compatibility type of the relationship or the kind of relationship it is, you might have to delete SnapMirror relationships, Storage Virtual Machine (SVM) peer relationships, and possibly cluster peer relationships.

The following lists what you must do to prepare SnapMirror and SnapVault relationships before reverting a cluster:

- Load-sharing mirror relationships
  You must delete all load-sharing mirror relationships and load-sharing mirror destination volumes.
- Data protection mirror relationships
  You must delete all data protection mirror relationships that are compatible with Data ONTAP 8.2 because they are not compatible with Data ONTAP 8.1.
- SnapVault relationships
  You must delete all SnapVault relationships because they do not exist prior to Data ONTAP 8.2.
- SVM peer relationships
  You must delete all SVM peer relationships because they do not exist prior to Data ONTAP 8.2.
- Cluster peer relationships
  After deleting all the previously listed relationships, you can try to revert the operating system without deleting cluster peer relationships. If any cluster peer relationships do need to be deleted, you are notified during the revert process.

Deleting load-sharing and data protection mirror relationships before reverting

Before reverting to Data ONTAP 8.1.x, you must delete all load-sharing mirror relationships and any data protection mirror relationships that were created in Data ONTAP 8.2.

Steps

1. View and save information about all of your mirror relationships and each relationship's destination volume by using the `snapmirror show` command and `volume show` commands, respectively.

   The following mirror relationships must be deleted:

   - All load-sharing mirror relationships
   - Any data protection mirror relationships that were created in Data ONTAP 8.2
     When viewing mirror relationships by using the `snapmirror show` command, these relationships have a `relationship capability` of `8.2 and above`.

- All data protection relationships if the cluster was re-created in Data ONTAP 8.2
  If the cluster was re-created in Data ONTAP 8.2 by using the `system configuration recovery cluster recreate` command, then all data protection relationships must be deleted.

Example

This example identifies the relationship capability of the mirror relationships in the cluster:

```
cluster1::> snapmirror show -fields relationship-capability
source-path destination-path relationship-capability
----------- ---------------- -----------------------
vs1:vol1    vs1:vol3         "8.2 and above"
```

2. Break each data protection relationship that you identified by using the `snapmirror break` command.

Example

```
cluster1::> snapmirror break -destination-path vs1:vol3
```

3. Delete each data protection mirror and load-sharing mirror relationship that you identified by using the `snapmirror delete` command.

Example

```
cluster1::> snapmirror delete -destination-path vs1:vol3
```

Only the mirror relationships are deleted; the destination volumes remain.

4. For each data protection mirror relationship that you deleted, use the `snapmirror release` command from the source node to remove the configuration information and snapshot copies from the source volume.

Example

```
cluster1::> snapmirror release -relationship-info-only -source-path vs1:vol1 -destination-
path vs1:vol3
```

5. Take each load-sharing destination volume offline by using the `volume offline` command.
6. Delete each load-sharing and uninitialized data protection destination volume by using the `volume delete` command.

After you finish

If desired, re-create the mirror relationships in the target reversion release.

Deleting data protection mirror relationships before reverting

Before reverting to Data ONTAP 8.1.x, you must break and delete any data protection mirror relationships that has the relationship capability of “8.2 and above” and you must break data protection mirror relationships with a relationship capability of “Pre-8.2.”

About this task

The extent of what you must do with data protection mirror relationships before reverting the cluster depends on whether the data protection mirror relationships have a relationship capability of “8.2 and above”:

- If the relationships have a relationship capability of “Pre-8.2,” you must release source volumes for relationships whose source volumes reside on the cluster you are reverting, and you must break relationships whose destination volumes reside on the cluster you are reverting.

  You do not need to do anything on cluster peers that have either source or destination volumes for the data protection mirror relationships you broke or released.
- If the relationships have a relationship capability of “8.2 or above,” you must remove the data protection mirror relationship entirely by releasing source volumes and by breaking and deleting destination volumes on the cluster you are reverting and on cluster peers.

The following procedure performs the tasks on the cluster you are reverting first, then performs tasks on the cluster peers. There are three clusters configured with a data protection mirror between cluster peers “cluster1” and “cluster2”, and a data protection mirror between cluster peers “cluster2” and “cluster3”. The procedure uses examples that describe the process for data protection mirror relationships that have a relationship capability of “8.2 or above.”

Steps

1. On the cluster you are reverting, remove data protection mirror relationships whose destination volumes are on the cluster and release source volumes on the cluster.

   a) On the cluster you are reverting, view information about the relationship capability of your mirror relationships by using the `snapmirror show` command with the `-fields relationship-capability` parameter.

   #### Example

   This example identifies the relationship capability of the mirror relationships in the cluster:

   ```
   cluster2::> snapmirror show -fields relationship-capability
   source-path destination-path relationship-capability
   ----------- ---------------- -----------------------
   cl1vs1:vol1 cl2vs2:vol2      "8.2 and above"
   ```

   b) Break each identified data protection mirror relationship by using the `snapmirror break` command.

Example

```
cluster2::> snapmirror break -destination-path cl2vs2:vol2
```

c) Delete each identified data protection mirror relationship by using the `snapmirror delete` command.

Example

```
cluster2::> snapmirror delete -destination-path cl2vs2:vol2
```

Only the mirror relationships are deleted; the destination volumes remain.

d) List any relationships that have source volumes on the cluster you are reverting by using the `snapmirror list-destinations` command.

Example

```
ie3070::> snapmirror list-destinations
                                                  Progress
Source             Destination         Transfer   Last
Relationship
Path         Type  Path         Status Progress   Updated      Id
----------- ----- ------------ ------- --------- ------------
---------------
cl2vs2:vol2
            DP    cl3vs3:vol3  -       -         -
590f2d55-fb9e-11e2-b2fc-123478912345
```

e) Remove the configuration information and Snapshot copies from the data protection mirror relationship whose source is on the cluster you are reverting by using the `snapmirror release` command.

Example

```
cluster2::> snapmirror release -relationship-info-only -source-path cl2vs2:vol2
-destination-path cl3vs3:vol3
```

2. On cluster peers that have either source volumes or destination volumes that are endpoints of relationships you previously broke and deleted or released, complete the removal of the relationships by removing or releasing the relationships on the cluster peers.

a) Using the destination volume information from step 1d, list the relationship compatibility by using the `snapmirror show` command with the `-fields relationship-compatibilty` parameter.

Example

```
cluster3::> snapmirror show -fields relationship-capability
source-path destination-path relationship-capability
----------- ---------------- -----------------------
```

```
cl2vs2:vol2 cl3vs3:vol3      "8.2 and above"
```

b) Break the relationship shown by the `snapmirror show` command by using the `snapmirror break` command on the cluster that contains the destination volume.

Example

```
cluster3::> snapmirror break -destination-path cl3vs3:vol3
```

c) Delete the relationship shown by the `snapmirror show` command by using the `snapmirror delete` command on the cluster that contains the destination volume.

Example

```
cluster3::> snapmirror delete -destination-path cl3vs3:vol3
```

d) Remove the configuration information and Snapshot copies by using the `snapmirror release` command on the cluster that contains the source volume.

This relationship is shown by the `snapmirror show` command of Step 1a.

Example

```
cluster1::> snapmirror release -relationship-info-only -source-path cl1vs1:vol1
-destination-path cl2vs2:vol2
```

After you finish

If desired, re-create the mirror relationships in the target reversion release.

Deleting SnapVault relationships before reverting

Before reverting to Data ONTAP 8.1.x, you must delete all SnapVault relationships that were created in Data ONTAP 8.2. The ability to back up a volume using SnapVault technology did not exist before Data ONTAP 8.2.

Steps

1. View and save information about all of your SnapVault relationships by using the `snapmirror show` command.

Example

```
cluster1::> snapmirror show
                                                                       Progress
Source            Destination  Mirror  Relationship  Total             Last
Path        Type  Path         State   Status        Progress  Healthy Updated
----------- ---- ------------ ------- -------------- --------- ------- --------
vs1:vol2    XDP  vs1:vol2_backup
```

```
                              Snapmirrored
                                      Idle           -         true    -
```

2. Change the privilege level of the command session to `advanced` by using the `set` command with the `-privilege advanced` parameter.
3. Break each SnapVault relationship that you identified by using the `snapmirror break` command with the `-delete-snapshots` parameter.

   Using the `-delete-snapshots` parameter deletes all of the Snapshot copies on the volume, a prerequisite for reverting if you have SnapVault relationships.

   Example

   ```
   cluster1::*> snapmirror break -destination-path vs1:vol2_backup
   -delete-snapshots
   ```

4. Delete each SnapVault relationship that you identified by using the `snapmirror delete` command.

   Example

   ```
   cluster1::*> snapmirror delete -destination-path vs1:vol2_backup
   ```

   Only the SnapVault relationships are deleted; the destination volumes remain.
5. For each SnapVault relationship that you deleted, use the `snapmirror release` command from the source volume to remove the configuration information and Snapshot copies from the source volume.

   Example

   ```
   cluster1::*> snapmirror release -source-path vs1:vol2 -destination-path vs1:vol2_backup
   ```

6. Take each destination volume offline by using the `volume offline` command.
7. Delete each destination volume by using the `volume delete` command.

Deleting SVM peer relationships before reverting

Before reverting to Data ONTAP 8.1.x, you must delete Storage Virtual Machine (SVM) peer relationships created in Data ONTAP 8.2 because they did not exist in Data ONTAP 8.1.x.

Before you begin

SnapMirror relationships must be deleted. The topic “Deleting data protection mirror relationships before reverting”contains detailed instructions.

Steps

1. Ensure that there are no SnapMirror relationships involving the cluster you are reverting by using the `snapmirror show` command.

   #### Example

   ```
   cluster2::> snapmirror show
   This table is currently empty.
   ```

2. To determine the SVM peer relationships you want to delete, view the SVM peer relationships on the cluster you are reverting by using the `vserver peer show` command.

   #### Example

   ```
   cluster2::> vserver peer show
                      Peer        Peer
          Vserver     Vserver     State
          ----------- ----------- ------------
          cl1vs1      cl2vs2
   peered
          cl2vs2      cl3vs3
   peered

          3 entries were displayed.
   ```

3. Delete each SVM peer relationship by using the `vserver peer delete` command.

   #### Example

   ```
   cluster2::> vserver peer delete -vserver cl1vs1
   -peer-vserver cl2vs2
   cluster2::> vserver peer delete -vserver cl2vs2
   -peer-vserver cl3vs3
   ```

Deleting cluster peer relationships before reverting

Before reverting to Data ONTAP 8.1.x, you might have to delete cluster peer relationships created or modified in Data ONTAP 8.2 because they are not compatible with the Data ONTAP 8.1.x format.

Before you begin

- All SnapMirror relationships must be deleted.
- All Storage Virtual Machine (SVM) peer relationships must be deleted.

About this task

Typically, you do not have to delete cluster peer relationships when reverting, but there are some cases in which you do have to delete cluster peer relationships. If deleting cluster peer relationships is required, a Data ONTAP 8.2 message notifies you to delete all SnapMirror relationships and cluster peer relationships after you initiate the cluster revert process by using the `system node revert-to` command.

Steps

1. Ensure that there are no SnapMirror relationships involving the cluster you are reverting by using the `snapmirror show` command.

   #### Example

   This example shows that there are no SnapMirror relationships.

   ```
   cluster2::> snapmirror show
   This table is currently empty.
   ```

2. Ensure that there are no SVM peer relationships configured on the cluster you are reverting by using the `vserver peer show` command.

   #### Example

   ```
   cluster2::> vserver peer show
   There are no Vserver peer relationships.
   ```

3. Show all of the configured cluster peer relationships by using the `cluster peer show` command.

   #### Example

   ```
   cluster2::> cluster peer show                                                   |
   Peer Cluster Name          Cluster Serial Number Availability                   |
   -------------------------- --------------------- ---------------                |
   cluster1                   1-80-123456           Available
   ```

4. Delete each cluster peer relationship by using the `cluster peer delete` command from the source node to remove the configuration information and Snapshot copies from the source volume.

   You must delete the cluster peer relationships from both clusters in the relationship.

   #### Example

   Assuming that cluster 1 and cluster 2 are in a cluster peer relationship, you delete the cluster peer relationship by using the following two commands:

   ```
   cluster1::> cluster peer delete -cluster cluster2
   cluster2::> cluster peer delete -cluster cluster1
   ```

SMB 3.0 must be disabled before reverting to a release earlier than Data ONTAP 8.2

Data ONTAP 8.2 and later releases provide support for SMB 3.0. Before you revert to a release earlier than Data ONTAP 8.2, you must disable SMB 3.0 on all Vservers in the cluster.

- The following advanced-level command disables this feature on all Vservers:

```
vserver cifs options modify -vserver * -smb3-enabled false
```

- You need to be aware that several new features and functionality in Data ONTAP 8.2 and later releases are dependent on SMB 3.0, including the following:
  - Continuously available share property
  - Persistent handles
  - Remote VSS for SMB shares
  - Witness
  - ODX copy offload
  - BranchCache version 2

  Note: The Hyper-V over SMB solution depends on the functionality provided by the continuously available share property, persistent handles, Remote VSS for SMB shares, and the Witness protocol. You should not disable SMB 3.0 until you have taken appropriate actions for any Hyper-V over SMB solutions on the cluster. Before disabling SMB 3.0, ensure that there are no ongoing Remote VSS shadow copy operations.

For more information, see the *Clustered Data ONTAP File Access and Protocols Management Guide*.

IPv6 configurations for CIFS must be reconfigured before reverting

Data ONTAP 8.2 and later releases provide support for IPv6 for CIFS. Before you revert to a release earlier than Data ONTAP 8.2, any configuration with only IPv6 addresses must be reconfigured with an IPv4 address.

This includes, but is not limited to the following:

- Vserver data LIFs, routing groups, and static routes for Vservers hosting CIFS servers
- Preferred domain controllers

Note: Domain controllers and other servers providing services to CIFS servers must have IPv4 addresses configured to allow network connectivity after a revert to a version of Data ONTAP that does not support IPv6 for CIFS.

For more information about managing CIFS server configurations, see the *Clustered Data ONTAP File Access and Protocols Management Guide*. For more information about configuring IPv4 and IPv6, see the *Clustered Data ONTAP Network Management Guide*.

Actions you must take for FPolicy configurations before reverting to a release earlier than Data ONTAP 8.2

Data ONTAP 8.2 and later releases provide support for the FPolicy feature. Before you revert to a release earlier than Data ONTAP 8.2, you must understand and act on some important revert considerations.

Before reverting to a version of Data ONTAP that does not support FPolicy, the following conditions must be met:

- Every file on which FPolicy servers set the offline bit must be either deleted or replaced with the original files before disabling FPolicy and reverting to a version of Data ONTAP that does not support FPolicy.

  Note: If you do not replace the files with the offline bit set with the original files prior to reverting, clients access the stub files instead of the files to which the stub refers.
- FPolicy functionality must be disabled on the cluster by disabling every FPolicy policy on the cluster.

For more information about managing FPolicy configurations, see the *Clustered Data ONTAP File Access and Protocols Management Guide*.

Actions you must take for Hyper-V over SMB solutions before reverting to a release earlier than Data ONTAP 8.2

Data ONTAP 8.2 and later releases provide support for the Hyper-V over SMB solution. Before you revert to a release earlier than Data ONTAP 8.2, you must understand and act on some important revert considerations.

You must consider the following and take action where necessary:

- There must be no file access by the Hyper-V servers to virtual machine files when you revert:
  - You can use the Hyper-V application to migrate virtual machine files to another storage device or to local storage.
  - You can power down all virtual machines and manually terminate Hyper-V server connections to the data LIFs.

    Data ONTAP disables SMB 3.0 before reverting; therefore, if the SMB connections are not manually terminated, Data ONTAP terminates them during the revert.
- You cannot use the Hyper-V over SMB solution if you revert to a version of Data ONTAP that does not support it.

  If you want to store virtual machine files on Vserver volumes, you must configure the Hyper-V servers to use connected LUNs to store and access virtual machine files. You must then copy the virtual machine files from the SMB shares to the connected LUNs.
- To revert, there can be no ongoing Remote VSS shadow copy operations.

  If there are any, you must wait for the operations to finish or manually abort them before proceeding with the revert. If you need to abort any shadow copy operations, contact technical support for assistance. Upon a revert, Data ONTAP does not delete existing Snapshot copies.
- You must disable the Shadow Copy feature on all Vservers before reverting.

  The following advanced privilege command disables this feature on all Vservers:

  ```
  vserver cifs options modify -vserver * -shadowcopy-enabled false
  ```
- You must disable SMB 3.0 on all Vservers before reverting.

  The following advanced privilege command disables this feature on all Vservers:

  ```
  vserver cifs options modify -vserver * -smb3-enabled false
  ```

  Note: Other CIFS functionality available on Data ONTAP 8.2 and later depends on SMB 3.0. You should not disable SMB 3.0 until you have taken appropriate action for other features and functionality that depends on SMB 3.0. This includes the continuously available share property, persistent handles, Remote VSS for SMB shares, the Witness protocol, ODX copy offload, the CIFS reparse point option, and BranchCache 2.

For more information about managing the Hyper-V over SMB solution, see the *Clustered Data ONTAP File Access and Protocols Management Guide*.

Actions you must take for local user and group configurations before reverting to a release earlier than Data ONTAP 8.2

Data ONTAP 8.2 and later releases provide support for local users and groups for CIFS servers. Before you revert to a release earlier than Data ONTAP 8.2, you must understand and act on some important revert considerations.

Before reverting to a version of Data ONTAP that does not support local users and groups for CIFS servers, you must consider the following and perform any applicable actions:

- Upon a revert to a prior major version of Data ONTAP, Data ONTAP does not use local users and groups during authentication and access token creation.

  If there are users who authenticate by using local user accounts and they need continued access after revert, you must configure domain accounts that those users can use to authenticate.
- Local users and groups are not removed from file and folder ACLs.

  File access requests that depend on access being granted because of permissions granted to local users or groups are denied. To allow access, you must reconfigure file permissions to allow access based on domain users and groups instead of local users and groups.
- You must disable the local users and groups feature on all Vservers before reverting.

  The following advanced privilege command disables this feature on all Vservers:

  ```
  vserver cifs options modify -vserver * -is-local-users-and-groups-
  enabled false
  ```

For more information about managing local users and groups and managing SMB access, see the *Clustered Data ONTAP File Access and Protocols Management Guide*.

Export policies must be configured for SMB access before reverting

Before you revert to a release earlier than Data ONTAP 8.2, you must configure export policies for SMB access on Vservers where data is accessed over SMB shares.

For releases earlier than Data ONTAP 8.2, export policies are mandatory for SMB access. Starting with Data ONTAP 8.2, export policies for SMB access are optional and are disabled by default. If you revert to a version of Data ONTAP where export policies are mandatory, export policies are turned on and required for SMB access. If you do not have export policies in place that permit SMB access prior to reverting, SMB clients are denied access to data over SMB shares.

The recommendation is that you configure export policies for SMB on all Vservers with SMB shares before you revert so that there are not hard-to-resolve SMB client access issues after the revert is complete.

For more information about configuring export policies for SMB access, see the *Clustered Data ONTAP File Access and Protocols Management Guide*.

BranchCache must be disabled before reverting to a release earlier than Data ONTAP 8.2

Data ONTAP 8.2 and later releases provide support for BranchCache. Before you revert to a release earlier than Data ONTAP 8.2, you must disable BranchCache on all Vservers in the cluster.

- The following command disables this feature on all Vservers:

  ```
  vserver cifs branchcache delete -vserver * -flush-hashes true
  ```

  This command deletes the BranchCache configuration and hashes on all Vservers. After the revert is complete, the storage administrator can manually delete the directories that contained hash stores, if desired.
- When you revert to a version of Data ONTAP that does not support BranchCache, the SMB shares do not advertise BranchCache capabilities to BranchCache-enabled clients; therefore, the clients do not request hash information.

  Instead, they request the actual content using normal SMB read requests.

For more information about managing BranchCache configurations, see the *Clustered Data ONTAP File Access and Protocols Management Guide*.

Reverting systems with deduplicated volumes

Before reverting from the Data ONTAP 8.2 release family, you must ensure that the volumes contain sufficient free space for the revert operation.

About this task

Reverting from a Data ONTAP 8.2 release family on a system that has deduplication enabled, includes running advanced mode commands. You must contact technical support for assistance.

If you have enabled both deduplication and data compression on a volume that you want to revert, then you must revert data compression before reverting deduplication.

Steps

1. Use the `volume efficiency show` command with the `-fields` option to view the progress of efficiency operations running on the volumes.

   #### Example

   The following command displays the progress of efficiency operations:

   ```
   volume efficiency show -fields vserver,volume,progress
   ```

2. Use the `volume efficiency stop` command with the `-all` option to stop all active and queued deduplication operations.

Example

The following command stops all active and queued deduplication operations on volume VolA:

```
volume efficiency stop -vserver vs1 -volume VolA -all
```

3. Use the `volume efficiency off` command to disable the deduplication operation.

Example

The following command disables deduplication operation on volume VolA:

```
volume efficiency off -vserver vs1 -volume VolA
```

4. Use the `set -privilege advanced` command to log in at the advanced privilege level.
5. Use the `volume efficiency revert-to` command with the `-version` option to downgrade the efficiency metadata of a volume to a specific version of Data ONTAP.

Example

The following command reverts the efficiency metadata on volume VolA to 8.1 version:

```
volume efficiency revert-to -vserver vs1 -volume VolA -version 8.1
```

> **Note:** The `volume efficiency revert-to` command reverts volumes that are present on the node on which this command is executed. This command does not revert volumes across nodes.

6. After the revert operation is complete, return to the admin privilege level by entering the following command:

```
set -privilege admin
```

For more information about deduplication, see the *Clustered Data ONTAP Logical Storage Management Guide*.

Reverting systems with compressed volumes

Before reverting from the Data ONTAP 8.2 release family, you must ensure that the volumes contain sufficient free space for the revert operation.

Before you begin

If you are reverting from Data ONTAP 8.2 release family, then you must have deleted all the Snapshot copies that have compressed data and decompress the data in the volumes.

About this task

Reverting from a Data ONTAP 8.2 release on a system that has data compression enabled, includes running advanced mode commands. You must contact technical support for assistance.

Steps

1. Use the `volume efficiency show` command with the `-fields` option to view the progress of efficiency operations running on the volumes.

   #### Example

   The following command displays the progress of efficiency operations:

   ```
   volume efficiency show -fields vserver,volume,progress
   ```

2. Use the `volume efficiency stop` command with the `-all` option to stop all active and queued data compression operations on the volume.

   #### Example

   The following command aborts both active and queued data compression operations on volume VolA:

   ```
   volume efficiency stop -vserver vs1 -volume VolA -all
   ```

3. Use the `volume efficiency modify` command to disable data compression on the volume.

   #### Example

   The following command disables data compression on volume VolA:

   ```
   volume efficiency modify -vserver vs1 -volume VolA -compression false -
   inline-compression false
   ```

4. Use the `volume efficiency off` command to disable deduplication on the volume.

   #### Example

   The following command disables deduplication on volume VolA:

   ```
   volume efficiency off -vserver vs1 -volume VolA
   ```

5. Use the `set -privilege advanced` command to log in at the advanced privilege level.

6. Use the `volume efficiency revert-to` command with the `-version` option to downgrade the efficiency metadata of a volume a Data ONTAP 8.1 release family.

   #### Example

   If you are reverting to Data ONTAP 8.1 release, the following command downgrades the efficiency metadata on volume VolA to 8.1 version:

   ```
   volume efficiency revert-to -vserver vs1 -volume VolA -version 8.1
   ```

   Note: The `volume efficiency revert-to` command reverts volumes that are present on the node on which this command is executed. This command does not revert volumes across nodes.

7. After the revert operation is complete, use the `set -privilege admin` command to return to the admin privilege.

   For more information about data compression, see the *Clustered Data ONTAP Logical Storage Management Guide*.

Deleting an SVM with Infinite Volume

If the cluster contains an Infinite Volume, you must delete the Infinite Volume and its containing Storage Virtual Machine (SVM) before reverting or downgrading from Data ONTAP 8.1.1 or a later release of the Data ONTAP 8.1 family to the first release in the Data ONTAP 8.1 family.

Before you begin

- The Infinite Volume must be unmounted.
- Any SnapMirror relationships with the Infinite Volume must be deleted.

Steps

1. If the Infinite Volume is online, use the `volume offline` command to take the Infinite Volume offline.
2. Use the `volume delete` command to delete the Infinite Volume.
3. Use the `volume delete` command to delete any FlexVol volumes that were created for data protection or recovery purposes.
4. Delete any customized user accounts and roles associated with the SVM.

   For information about users and roles, see the *Clustered Data ONTAP File Access Management Guide for CIFS* and the *Clustered Data ONTAP File Access Management Guide for NFS*.
5. Use the `vserver stop` command to stop the SVM.
6. Use the `vserver delete` command to delete the SVM.
7. Reboot all the nodes in the cluster.

Verifying the system time

You should verify that NTP is configured, and that the time is synchronized across the cluster.

About this task

For more information about managing the system time, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

Steps

1. Use the `system services ntp server show` command to verify that each node is associated with an NTP server.

Example

```
cluster1::> system services ntp server show
Node   Server               Version
------ -------------------- -----------
node0
       ntp1.example.com     max
       ntp2.example.com     max
node1
       ntp1.example.com     max
       ntp2.example.com     max
node2
       ntp1.example.com     max
       ntp2.example.com     max
node3
       ntp1.example.com     max
       ntp2.example.com     max
```

2. Verify that each node has the same date and time:

If you are running Data ONTAP...	Enter the following command...
8.1.x	**`system node date show`**
8.2.x	**`cluster date show`**

Example

```
cluster1::> cluster date show
Node      Date                Timezone
--------- ------------------- -------------------------
node0     4/6/2013 20:54:38   GMT
node1     4/6/2013 20:54:38   GMT
node2     4/6/2013 20:54:38   GMT
node3     4/6/2013 20:54:38   GMT
4 entries were displayed.
```

Preparing Snapshot copies before reverting

Before reverting to an earlier Data ONTAP release, you must delete any Snapshot copies that were created after upgrading to the current release, delete root aggregate and root volume Snapshot copies, and disable Snapshot schedules for the root aggregate and root volume.

Before you begin

If you are reverting in a SnapMirror environment, you must first have deleted the following mirror relationships:

- All load-sharing mirror relationships
- Any data protection mirror relationships that were created in Data ONTAP 8.2.x
- All data protection mirror relationships if the cluster was re-created in Data ONTAP 8.2.x

FlexClone volumes that contain Snapshot copies that were created in Data ONTAP 8.2 or later must be either split from their parents or deleted. You will not be able to delete the Snapshot copies until these FlexClone volumes are split or deleted.

About this task

Root aggregate and root volume Snapshot copies are not visible from the clustershell and are not deleted by a `snapshot delete` command; therefore, you must delete the Snapshot copies manually.

Steps

1. Identify any Snapshot copies that were created after upgrading to the current release:

   `volume snapshot show -fs-version 8.2`

2. Delete the Snapshot copies that you identified:

   `volume snapshot delete {-fs-version 8.2 -node` ***`nodename`*** **`}`**

3. Identify the root aggregate on each node in the cluster by using the `run -node` *`nodename`* `aggr status` command.

   You identify the root aggregate by the word root in the options column of the `aggr status` command output.

   #### Example

   ```
   cluster1::> run -node node1 aggr status

              Aggr State           Status            Options
             aggr0 online          raid_dp, aggr     root
                                   64-bit
   ```

4. Disable Snapshot schedules on the root aggregate:

   `run -node` ***`nodename`*** **`aggr options`** ***`root_aggr_name`*** **`nosnap on`**

5. Delete the root aggregate Snapshot copies:

   `run -node` ***`nodename`*** **`snap delete -A -a -f`** ***`aggr_name`***

6. Identify the root volume on each node in the cluster by using the `run -node` *`nodename`* `vol status` command.

   You identify the root volume by the word root in the Options column of the `vol status` command output.

   #### Example

   ```
   vs1::>  run -node node1 vol status
   ```

```
          Volume State            Status            Options
            vol0 online           raid_dp, flex     root, nvfail=on
                                  64-bit
```

7. Disable Snapshot schedules on the root volume:

   run -node *nodename* vol options *root_volume_name* nosnap on

8. Delete the root volume Snapshot copies:

   run -node *nodename* snap delete -a -f *volume_name*

Obtaining Data ONTAP software images

You must copy a software image from the N series support website (accessed and navigated as described in *Websites*) to an HTTP server on your network so that nodes can access the images by using the `system node image update` command.

About this task

To upgrade, revert, or downgrade the cluster to the target release of Data ONTAP, you need access to software images. Software images, firmware version information, and the latest firmware for your platform model are available on the N series support website (accessed and navigated as described in *Websites* on page 7). Note the following important information:

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

Steps

1. Locate the target Data ONTAP software on the N series support website (accessed and navigated as described in *Websites* on page 7).
2. Copy the software image (for example, `821_q_image.tgz`) from the N series support website (accessed and navigated as described in *Websites* on page 7) to the directory on the HTTP server from which the image will be served.

Installing Data ONTAP software images in a cluster

If needed, install the target Data ONTAP 8.x image software package, but leave the default set to the current Data ONTAP 8.x version.

Before you begin

You must have obtained the Data ONTAP software images.

Steps

1. Choose one of the following options based on your requirements:

If you want to...	Enter this command...
Download, but not install, the software image	**system node image get -node * -package *location* -replace-package true -background true** This command downloads the software image to all of the nodes simultaneously. To download the image to each node one at a time, do not specify the `-background` parameter.
Install a previously downloaded software image	**system node image update -node * -package file:///mroot/etc/software/*image_name* -background true** Note the following considerations for this command: • If you are currently running Data ONTAP 8.2, the `-package` parameter only requires you to enter the image name; you do not need to enter the full directory path. • If you are unsure of the image name to install, you can view a list of previously downloaded software images by using the `system node image package show` command. • This command installs the software image on all of the nodes simultaneously. To install the image on each node one at a time, do not specify the `-background` parameter.
Download and install the software image in the same operation	**system node image update -node * -package *location* -replace-package true -background true** This command downloads and installs the software image on all of the nodes simultaneously. To download and install the image on each node one at a time, do not specify the `-background` parameter.

2. Verify that the software image is downloaded and installed on each node:
system node image show-update-progress -node *

This command displays the current status of the software image download and installation. You should continue to run this command until all nodes report a `Run Status` of `Exited`, and an `Exit Status` of `Success`.

Example

This example shows a 2-node cluster in which the software image has been downloaded and installed successfully on both nodes:

```
cluster1::> system node image show-update-progress -node *
There is no update/install in progress
Status of most recent operation:
        Run Status:     Exited
        Exit Status:    Success
        Phase:          Run Script
        Exit Message:   Installation complete. image2 updated on node node0.
There is no update/install in progress
Status of most recent operation:
        Run Status:     Exited
```

```
         Exit Status:    Success
         Phase:          Run Script
         Exit Message:   Installation complete. image2 updated on node node1.
2 entries were acted on.
```

Reverting a Data ONTAP cluster

To take the cluster offline to revert to an earlier Data ONTAP release, you must disable storage failover and the data LIFs, address reversion preconditions, revert the cluster and filesystem configurations on a node, and then repeat the process for each additional node in the cluster.

Before you begin

You must have satisfied reversion preparation requirements.

About this task

Reverting a cluster requires you to take the cluster offline for the duration of the reversion.

Steps

1. Verify that the target Data ONTAP software is installed:

   `system node image show`

   #### Example

   This example shows that version 8.1.2 is installed as the alternate image on both nodes.

   ```
   cluster1::> system node image show
                    Is      Is                 Install
   Node     Image   Default Current Version    Date
   -------- ------- ------- ------- --------   -------------------
   node0
            image1  true    true    8.2.1      11/25/2013 12:37:36
            image2  false   false   8.1.2      4/22/2013 13:52:22
   node1
            image1  true    true    8.2.1      11/25/2013 12:41:16
            image2  false   false   8.1.2      4/22/2013 13:55:22
   4 entries were displayed.
   ```

   For more information about installing the target Data ONTAP software image, see *Installing Data ONTAP 8.x software images in a cluster*.

2. Set the target Data ONTAP software image to be the default image:

   `system image modify {-node * -iscurrent false} -isdefault true`

3. Verify that the target Data ONTAP software image is set as the default image:

   `system node image show`

Example

This example shows that version 8.1.2 is set as the default image on both nodes.

```
cluster1::> system node image show
                 Is      Is                 Install
Node     Image   Default Current Version    Date
-------- ------- ------- ------- --------   -------------------
node0
         image1  false   true    8.2.1      11/25/2013 12:37:36
         image2  true    false   8.1.2      4/22/2013 13:52:22
node1
         image1  false   true    8.2.1      11/25/2013 12:41:16
         image2  true    false   8.1.2      4/22/2013 13:55:22
4 entries were displayed.
```

4. Disable all of the data LIFs in the cluster:

`network interface modify {-role data} -status-admin down`

5. If the cluster consists of only two nodes, disable cluster HA:

`cluster ha modify -configured false`

6. Disable storage failover for the nodes in the HA pair by entering the following command for either node:

`storage failover modify -node` ***`nodename`*** **`-enabled false`**

You only need to disable storage failover once for the HA pair. When you disable storage failover for a node, storage failover is also disabled on the node's partner.

7. Set the privilege level to advanced:

`set -privilege advanced`

8. If the cluster consists of only two nodes, then verify that the node does not hold epsilon by doing the following:

a) Check to see if the node currently holds epsilon:

`cluster show -node` ***`nodename`***

Example

This example shows that the node holds epsilon.

```
cluster1::*> cluster show -node node1

          Node: node1
          UUID: 026efc12-ac1a-11e0-80ed-0f7eba8fc313
       Epsilon: true
   Eligibility: true
        Health: true
```

b) If the node holds epsilon, then mark epsilon false on the node so that epsilon can be transferred to the node's partner:

`cluster modify -node` ***`nodenameA`*** **`-epsilon false`**

c) Transfer epsilon to the node's partner by marking epsilon true on the partner node:

```
cluster modify -node nodenameB -epsilon true
```

9. Verify that the node is ready for reversion:

   ```
   system node revert-to -node nodename -check-only true -version 8.1
   ```

   The `check-only` parameter identifies any preconditions that must be addressed before reverting, such as the following examples:

   - Disabling storage failover
   - Disabling the Snapshot policy
   - Deleting Snapshot copies that were created after upgrading to the later release family

   Proceed to the next step when all identified preconditions have been addressed.

10. Verify that all of the preconditions have been addressed:

    ```
    system node revert-to -node nodename -check-only true -version 8.1
    ```

11. Revert the cluster configuration of the node:

    ```
    system node revert-to -node nodename -version 8.1
    ```

    The `-version` option refers to the target release family. For example, if the software you installed and verified in Step 2 is Data ONTAP 8.1.2, the correct value of the `-version` option is `8.1`.

    The cluster configuration is reverted, and then you are logged out of the clustershell.

12. Log back into the clustershell, and then switch to the nodeshell:

    ```
    system node run -node nodename
    ```

13. Revert the filesystem configuration of the node:

    ```
    revert_to 8.1c
    ```

    This command verifies that the node's filesystem configuration is ready to be reverted and then reverts it. If any preconditions are identified, you must address them and then reenter the `revert_to` command.

    When the command finishes, the LOADER prompt is displayed.

14. At the LOADER prompt, enter the following command to boot to the target release:

    ```
    boot_ontap
    ```

15. Repeat Steps *8* on page 120 through *14* on page 121 on the other node in the HA pair.

16. If the cluster consists of only two nodes, reenable cluster HA:

    ```
    cluster ha modify -configured true
    ```

17. Reenable storage failover on both nodes if it was previously disabled:

    ```
    storage failover modify -node nodename -enabled true
    ```

18. Repeat Steps *6* on page 120 through *17* on page 121 for each additional HA pair in the cluster.

Completing post-reversion tasks

After reverting to an earlier Data ONTAP release family, you might need to perform additional tasks to ensure cluster health and storage availability.

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

Verifying that the cluster is in quorum

Before and after you perform an upgrade, reversion, or downgrade, you must ensure that all nodes are participating in a replicated database (RDB) quorum and that all rings are in the quorum. You must also verify that the per-ring quorum master is the same for all nodes.

About this task

For more information about cluster replication rings and RDB quorums, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

Steps

1. Set the privilege level to advanced:

   `set -privilege advanced`

   Enter **`y`** to continue.

2. Display each RDB process:

To display this RDB process...	Enter this command...
Management application	**`cluster ring show -unitname mgmt`**
Volume location database	**`cluster ring show -unitname vldb`**
Virtual-Interface manager	**`cluster ring show -unitname vifmgr`**
SAN management daemon	**`cluster ring show -unitname bcomd`**

Example

This example shows the volume location database process for a cluster running Data ONTAP 8.1.x:

```
cluster1::*> cluster ring show -unitname vldb

Node   UnitName Epoch DB Epoch DB Trnxs Master
------ -------- ----- -------- -------- ---------
node0  vldb     154   154      14847    node0
```

```
node1  vldb     154   154      14847    node0
node2  vldb     154   154      14847    node0
node3  vldb     154   154      14847    node0
4 entries were displayed.
```

Example

This example shows the volume location database process for a cluster running Data ONTAP 8.2.x:

```
cluster1::*> cluster ring show -unitname vldb
Node      UnitName Epoch    DB Epoch DB Trnxs Master    Online
--------- -------- -------- -------- -------- --------- ---------
node0     vldb     154      154      14847    node0     master
node1     vldb     154      154      14847    node0     secondary
node2     vldb     154      154      14847    node0     secondary
node3     vldb     154      154      14847    node0     secondary
4 entries were displayed.
```

For each process, verify the following configuration details:

- The relational database epoch and database epochs match for each node.
- The per-ring quorum master is the same for all nodes.
  Note that each ring might have a different quorum master.

3. Return to the admin privilege level:

`set -privilege admin`

4. If you are operating in a SAN environment, verify that each node is in a SAN quorum:

`event log show -messagename scsiblade.*`

The most recent `scsiblade` event message for each node should indicate that the scsi-blade is in quorum. During the upgrade, reversion, or downgrade process, each node will temporarily fall out of SAN quorum. Therefore, if you are verifying the SAN quorum after completing an upgrade, reversion, or downgrade, you may notice critical event messages warning you that the nodes were previously out of SAN quorum.

If a node is out of SAN quorum, you can use the `storage failover takeover` and `storage failover giveback` commands to perform a planned takeover and giveback with the node's high-availability partner and bring the node back into SAN quorum.

Example

In this example, both nodes in the cluster are in SAN quorum.

```
cluster1::> event log show -messagename scsiblade.*
Time                Node             Severity      Event
------------------- ---------------- ------------- --------------------------
8/13/2013 14:03:51  node0            INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
8/13/2013 14:03:51  node1            INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
```

Example

This example shows a two-node cluster after performing an upgrade. Each node shows a previous out of SAN quorum event message from when the node was upgraded. However, the most recent event message for each node shows that both nodes are presently in SAN quorum.

```
cluster1::> event log show -messagename scsiblade.*
Time                Node             Severity      Event
------------------- ---------------- ------------- ---------------------------
8/13/2013 15:37:51  node1            INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
8/13/2013 15:31:26  node1            CRITICAL      scsiblade.out.of.quorum: This node ...
8/13/2013 15:31:26  node1            INFORMATIONAL scsiblade.ics.trace.dumpfile: The ...
8/13/2013 15:31:26  node1            INFORMATIONAL scsiblade.ics.trace.dumpfile: The ...
8/13/2013 15:30:43  node0            INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
8/13/2013 15:24:16  node0            CRITICAL      scsiblade.out.of.quorum: This node ...
8/13/2013 15:24:16  node0            INFORMATIONAL scsiblade.ics.trace.dumpfile: The ...
8/13/2013 15:24:16  node0            INFORMATIONAL scsiblade.ics.trace.dumpfile: The ...
```

Verifying cluster and SVM health

Before and after you upgrade, revert, or downgrade a cluster, you should verify that the nodes are healthy and eligible to participate in the cluster, and the aggregates and volumes are online.

Steps

1. Verify that the nodes in the cluster are online and are eligible to participate in the cluster:

   `cluster show`

   #### Example

   ```
   cluster1::> cluster show
   Node                  Health  Eligibility
   --------------------- ------- ------------
   node0                 true    true
   node1                 true    true
   ```

   If any node is unhealthy or ineligible, check EMS logs for errors and take corrective action. For more information about EMS messages, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

2. Determine if any disk drives are broken, undergoing maintenance, or reconstructing:

To check for...	Do this...
Broken disks	**a.** Display any broken disks: **`storage disk show -state broken`** **b.** Remove or replace any broken disks.

To check for...	Do this...
Disks undergoing maintenance or reconstructing	**a.** Display any disks in maintenance, pending, or reconstructing states: **`storage disk show -state maintenance|pending|reconstructing`** **b.** Wait for the maintenance or reconstruction operation to complete before proceeding.

3. To verify that all aggregates are online, display the state of physical and logical storage, including storage aggregates:

`storage aggregate show -state !online`

This command displays the aggregates that are *not* online.

Example

```
cluster1::> storage aggregate show -state !online
There are no entries matching your query.
```

For more information about managing aggregates, see the *Clustered Data ONTAP Physical Storage Management Guide*.

4. To verify that all volumes are online, display any volumes *not* online:

`volume show -state !online`

Example

```
cluster1::> volume show -state !online
There are no entries matching your query.
```

For more information about managing volumes, see the *Clustered Data ONTAP Logical Storage Management Guide*.

Enabling and reverting LIFs to home ports

During a reboot, some LIFs might have been migrated to their assigned failover ports. Before and after you upgrade, revert, or downgrade a cluster, you must enable and revert any LIFs that are not on their home ports.

About this task

The `network interface revert` command reverts a LIF that is not currently on its home port back to its home port, provided that the home port is operational. A LIF's home port is specified when the LIF is created; you can determine the home port for a LIF by using the `network interface show` command.

Steps

1. Display the status of all LIFs:

   `network interface show`

   #### Example

   This example displays the status of all LIFs for a Storage Virtual Machine (SVM, formerly known as Vserver).

   ```
   cluster1::> network interface show -vserver vs0
               Logical    Status     Network            Current       Current Is
   Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
   ----------- ---------- ---------- ------------------ ------------- ------- ----
   vs0
               data001    down/down  192.0.2.120/24     node0         e0e     true
               data002    down/down  192.0.2.121/24     node0         e0f     true
               data003    down/down  192.0.2.122/24     node0         e2a     true
               data004    down/down  192.0.2.123/24     node0         e2b     true
               data005    down/down  192.0.2.124/24     node0         e0e     false
               data006    down/down  192.0.2.125/24     node0         e0f     false
               data007    down/down  192.0.2.126/24     node0         e2a     false
               data008    down/down  192.0.2.127/24     node0         e2b     false
   8 entries were displayed.
   ```

   If any LIFs appear with a `Status Admin` status of `down` or with an `Is home` status of `false`, continue with the next step.

2. Enable the data LIFs:

   `network interface modify {-role data} -status-admin up`

   #### Example

   ```
   cluster1::> network interface modify {-role data} -status-admin up
   8 entries were modified.
   ```

3. Revert LIFs to their home ports:

   `network interface revert *`

   #### Example

   This command reverts all LIFs back to their home ports and changes all LIF home statuses to `true`.

   ```
   cluster1::> network interface revert *
   8 entries were acted on.
   ```

4. Verify that all LIFs are in their home ports:

   `network interface show`

   #### Example

   This example shows that all LIFs for SVM vs0 are on their home ports.

```
cluster1::> network interface show -vserver vs0
            Logical    Status     Network            Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
----------- ---------- ---------- ------------------ ------------- ------- ----
vs0
            data001      up/up    192.0.2.120/24     node0         e0e     true
            data002      up/up    192.0.2.121/24     node0         e0f     true
            data003      up/up    192.0.2.122/24     node0         e2a     true
            data004      up/up    192.0.2.123/24     node0         e2b     true
            data005      up/up    192.0.2.124/24     node1         e0e     true
            data006      up/up    192.0.2.125/24     node1         e0f     true
            data007      up/up    192.0.2.126/24     node1         e2a     true
            data008      up/up    192.0.2.127/24     node1         e2b     true
8 entries were displayed.
```

Preparing Snapshot copies after reverting

After reverting to an earlier version of Data ONTAP, you must enable Snapshot schedules of the root volume and root aggregate to start creating Snapshot copies again.

About this task

You are reenabling the Snapshot schedules of the root volume and root aggregate that you disabled before you reverted to an earlier version of Data ONTAP.

Steps

1. Enable the Snapshot schedule of the root volume by using the `run -node` *`nodename`* `vol options` *`root_vol_name`* `nosnap off` command.

   Example

   ```
   cluster1::> run -node node1 vol options vol0 nosnap off
   ```

2. Enable the Snapshot schedule of the root aggregate by using the `run -node` *`nodename`* `aggr options` *`root_aggr_name`* `nosnap off` command.

   Example

   ```
   cluster1::> run -node node1 aggr options aggr0 nosnap off
   ```

Verifying client access (CIFS and NFS)

For the configured protocols, test access from CIFS and NFS clients to verify that the cluster is accessible.

Considerations for downgrading the SP firmware

If your current SP firmware version is not supported for the Data ONTAP release you are downgrading or reverting to, you must install a supported SP firmware version for the earlier Data ONTAP release.

After the Data ONTAP revert or downgrade process is complete, you must take the action to install an SP firmware version that is supported for the Data ONTAP version you reverted or downgraded to. For information about downloading and installing an SP firmware version, see the N series support website (accessed and navigated as described in *Websites* on page 7).

Reinstalling the required gateway license after revert or downgrade

The licensing scheme for gateways changed in Data ONTAP 8.2. Data ONTAP 8.2 and later requires the V_StorageAttach license package for gateways to be able to access LUNs on storage arrays. Depending on the type of reversion you do, you might need to install the gateway license key for a prior release.

A license key is required for each gateway being reverted or downgraded.

If...	Then the gateway license requirement is...
The gateway was upgraded to Data ONTAP 8.2 from the 8.1.x release family and you want to revert to a release in the 8.1.x release family again	You do not have to install the gateway license for the release to which the system is being reverted. Data ONTAP remembers the gateway license from when it was upgraded from 8.1.x to 8.2; Data ONTAP reinstalls those licenses when Data ONTAP is downgraded to 8.1.x.
Data ONTAP 8.2 is the first release of Data ONTAP installed on the system and you want to revert to a release in the 8.1.x release family	You must manually install the license key for gateways that is supported for the release to which the system is being reverted. Contact your sales representative to obtain the appropriate license key for your system and release.
You want to downgrade the system from Data ONTAP 8.2.x to an earlier release in the 8.2 release family	You do not have to install the V_StorageAttach license package; Data ONTAP remembers the appropriate license for the release family.

For information about how to install licenses, see the *Clustered Data ONTAP System Administration Guide for SVM Administrators*.

Enabling capability for SVM administrator with vsadmin role

After reverting to an earlier version of clustered Data ONTAP, the event generate-autosupport-log capability becomes unavailable for the predefined role vsadmin. Therefore, you must create a custom role with command directory name as `event generate autosupport log` and create a user for the new role.

About this task

With the new custom role, the SVM administrator can execute the `event generate-autosupport-log` command or zapi equivalent.

Steps

1. Use the `security login role create` command to create a new SVM administrator role.

   #### Example

   The following example shows how to create a new role test:

   ```
   cluster1::>security login role create -role test -cmddirname "event
   generate-autosupport-log" -access all -query "" -vserver vs1
   ```

2. Use the `security login create -username` command to create a new user and assign new role to the user.

   #### Example

   The following example shows how to create a new user and assign the new role:

   ```
   cluster1::>security login create -username user_test1 -application
   ontapi -authmethod password -role test -vserver vs1
   ```

Downgrading clusters to an earlier release in the same release family

Transitioning a cluster to an earlier release in the same Data ONTAP release family is referred to as a *downgrade*. Doing so requires preparation, downloading and booting the earlier release, and completing post-downgrade procedures.

Downgrading does not require modifications to Data ONTAP on-disk structures; you must simply 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 clusters, but you should call technical support if you encounter problems during or after upgrading, or if you want to downgrade a production cluster.

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 cluster and you want to return to the original release when testing is completed.
- You are configuring a new cluster—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 in a production environment without assistance. If you encounter any of the following circumstances, contact technical support immediately:

- The upgrade process fails and cannot finish.
- The upgrade process finishes, but the cluster is unusable in a production environment.
- The upgrade process finishes and the cluster goes into production, but you are not satisfied with its behavior.
- The upgrade process finishes for some but not all of the nodes, and you decide that you want to downgrade.

Planning your downgrade

Because new features are introduced in each release of Data ONTAP, you must understand downgrade requirements and evaluate how they might impact your current configuration.

Before proceeding with the downgrade, you should plan to do the following:

- Review the *Release Notes* for the Data ONTAP downgrade source release.

- Understand any requirements for downgrading to the target release from your existing software.
- Note any potential functionality changes to your cluster after the downgrade.
- Be prepared to address all points in the downgrade checklist.

Cluster downgrade checklist

You can use this checklist to record your progress as you prepare for the downgrade, perform the downgrade, and complete post-downgrade tasks.

Steps for preparing to downgrade

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

Condition	**Complete?**
Software and hardware support in the target release is confirmed. To confirm hardware support, visit *IBM System Storage N series Introduction and Planning Guide* at the N series support website (accessed and navigated as described in *Websites* on page 7)	
All release-specific downgrade issues have been resolved.	
You have clustershell access privileges.	
The cluster and Storage Virtual Machines (SVMs) are running and healthy. All aggregates and volumes should be healthy and online before proceeding with the reversion. You can use the `cluster show` command to verify the status of the nodes.	
The cluster is in quorum. All nodes are participating in a quorum and all rings are in the quorum. The per-ring quorum master should be the same for all nodes.	
LIFs are online and on their correct home ports. You can use the `network interface` command to display and modify LIF configuration.	
The system time is synchronized across the cluster.	
If you are preparing to downgrade a Gateway, no back-end configuration errors exist. You can check for back-end configuration errors using the `storage array config show` command.	
Any CIFS sessions that are not continuously available have been terminated.	

Condition	Complete?
Each node is running Data ONTAP 8.2.0 or later. Ensure that you are running on the minimum software version allowed for the downgrade by running the `system node image show` command.	
You have the target Data ONTAP software image available on an HTTP server. Download the software image for the target Data ONTAP release from the N series support website (accessed and navigated as described in *Websites* on page 7), and place it on an HTTP server that is accessible by each node.	
The target Data ONTAP software images are installed on each node and set as the alternate boot device image. You can use the `system node image update` command to install the software images. You can use the `system node image show` command to verify that the software images are installed as the alternate boot image on each node.	
SnapMirror operations are suspended.	

Steps for performing a nondisruptive downgrade

The nondisruptive downgrade is complete when all of the following steps have been completed:

Condition		Complete?
No jobs are running. If any aggregate, volume, mirror, NDMP (dump or restore), or Snapshot jobs are running or queued, allow the jobs to complete successfully or stop the queued entries.		
The target Data ONTAP software is installed and set as the default boot image.		
Each HA pair is downgraded.	The first node in the HA pair is downgraded.	
	The node's partner is downgraded.	

Steps for performing a disruptive downgrade

Condition	Complete?
No jobs are running. If any aggregate, volume, mirror, NDMP (dump or restore), or Snapshot jobs are running or queued, allow the jobs to complete successfully or stop the queued entries.	
The target Data ONTAP software is installed and set as the default boot image.	
Each node is rebooted.	

Steps for after downgrading

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

Condition	Complete?
The cluster is in quorum. Ensure that all nodes are participating in a quorum and all rings are in the quorum. Verify also that the per-ring quorum master is the same for all nodes.	
The cluster and SVMs are running and healthy. All aggregates and volumes should be healthy and online after the reversion. You can use the `cluster show` command to verify the status of the nodes.	
LIFs are online and on their correct home ports. You can use the `network interface` command to display and modify LIF configuration.	
Client access is verified. You should verify that clients can access the cluster for each configured protocol.	
SnapMirror operations are resumed.	
The Service Processor firmware version is confirmed. If your current SP firmware version is not supported for the Data ONTAP release to which you reverted, you must install a supported SP firmware version for the earlier Data ONTAP release.	

What to check before downgrading gateways

The process of downgrading systems with native disks is the same for gateways and filers. Gateways that use array LUNs require some special checks in addition to the checks for all systems running Data ONTAP.

Before downgrading a system running Data ONTAP 8.2 that uses array LUNs, check for the following:

- Whether you are using functionality that is supported in Data ONTAP 8.2 but not in the release to which you are downgrading your system
- Whether there are back-end storage related configuration errors

Storage Encryption downgrade restriction

If your storage system is configured for Storage Encryption, do not attempt to install and run any Data ONTAP version earlier than 8.2.1. Doing so would disable your system and render your data inaccessible.

Downgrade process considerations

Make sure that you know about downgrade issues and limitations before downgrading clusters to an earlier version of Data ONTAP.

Be aware of the following issues:

- After you upgrade clusters to Data ONTAP 8.2, you can downgrade to an earlier release in the Data ONTAP 8.2 release family.

  For example, you can downgrade from Data ONTAP 8.2.1 to 8.2.0.
- Downgrading affects all nodes in the cluster.

  The downgrade must be performed on all nodes in the cluster; however, some of the procedures must be performed on each HA pair and completed on each set of nodes before other pairs are downgraded.
- You can downgrade Data ONTAP nondisruptively.

  During the downgrade process, the cluster remains online and continues to serve data.
- If your cluster serves CIFS clients, nondisruptive downgrades are supported for Hyper-V over SMB solutions.

  Hyper-V over SMB solutions enable Hyper-V and the contained virtual machines to remain online and to provide continuous availability during the Data ONTAP downgrade. For more information, see the *Clustered Data ONTAP File Access Management Guide for CIFS*.

  For all other CIFS configurations, client sessions are terminated. You should direct users to end their sessions before you downgrade to prevent data loss.
- Data ONTAP clusters can operate for a limited time in a *mixed version* state, in which nodes in a cluster are running Data ONTAP versions from different release families. However, the upgrade is not complete until all nodes are running the new target release.

When the cluster is in a mixed-version state, you should not enter any commands that alter the cluster operation or configuration except as necessary to satisfy upgrade requirements; monitoring operations are permitted.

Identifying downgrade issues

Every Data ONTAP release family has unique downgrade requirements that you must understand and resolve before you decide to downgrade.

For additional information, and to check for downgrade issues that might have been discovered later, see the *Clustered Data ONTAP Release Notes*. The following list summarizes downgrade issues known when this guide was published:

- If your current SP firmware version is not supported for the Data ONTAP release that you are downgrading or reverting to, you must install a supported SP firmware version for the earlier Data ONTAP release.
- If you downgrade or revert to a release earlier than Data ONTAP 8.2.1, you must run the `security ssh prepare-to-downgrade` command at the advanced privilege level to reset the SSH security configurations of the cluster and all Storage Virtual Machines (SVMs) to the default settings used in the earlier release.
- Exporting qtrees is not supported in releases earlier than Data ONTAP 8.2.1.

  If you plan to downgrade or revert to an earlier release, you must first take certain actions; otherwise, the downgrade or revert operation fails.

Preparing for the Data ONTAP downgrade process

Before downgrading, you need to verify cluster health, verify aggregate and volume health, and install the target Data ONTAP image.

Verifying that the cluster is in quorum

Before and after you perform an upgrade, reversion, or downgrade, you must ensure that all nodes are participating in a replicated database (RDB) quorum and that all rings are in the quorum. You must also verify that the per-ring quorum master is the same for all nodes.

About this task

For more information about cluster replication rings and RDB quorums, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

Enter **`y`** to continue.

2. Display each RDB process:

To display this RDB process...	Enter this command...
Management application	**`cluster ring show -unitname mgmt`**
Volume location database	**`cluster ring show -unitname vldb`**
Virtual-Interface manager	**`cluster ring show -unitname vifmgr`**
SAN management daemon	**`cluster ring show -unitname bcomd`**

Example

This example shows the volume location database process for a cluster running Data ONTAP 8.1.x:

```
cluster1::*> cluster ring show -unitname vldb

Node   UnitName Epoch DB Epoch DB Trnxs Master
------ -------- ----- -------- -------- ---------
node0  vldb     154   154      14847    node0
node1  vldb     154   154      14847    node0
node2  vldb     154   154      14847    node0
node3  vldb     154   154      14847    node0
4 entries were displayed.
```

Example

This example shows the volume location database process for a cluster running Data ONTAP 8.2.x:

```
cluster1::*> cluster ring show -unitname vldb
Node      UnitName Epoch    DB Epoch DB Trnxs Master    Online
--------- -------- -------- -------- -------- --------- ---------
node0     vldb     154      154      14847    node0     master
node1     vldb     154      154      14847    node0     secondary
node2     vldb     154      154      14847    node0     secondary
node3     vldb     154      154      14847    node0     secondary
4 entries were displayed.
```

For each process, verify the following configuration details:

- The relational database epoch and database epochs match for each node.
- The per-ring quorum master is the same for all nodes.
  Note that each ring might have a different quorum master.

3. Return to the admin privilege level:

   `set -privilege admin`

4. If you are operating in a SAN environment, verify that each node is in a SAN quorum:

   `event log show -messagename scsiblade.*`

The most recent `scsiblade` event message for each node should indicate that the scsi-blade is in quorum. During the upgrade, reversion, or downgrade process, each node will temporarily fall out of SAN quorum. Therefore, if you are verifying the SAN quorum after completing an upgrade, reversion, or downgrade, you may notice critical event messages warning you that the nodes were previously out of SAN quorum.

If a node is out of SAN quorum, you can use the `storage failover takeover` and `storage failover giveback` commands to perform a planned takeover and giveback with the node's high-availability partner and bring the node back into SAN quorum.

Example

In this example, both nodes in the cluster are in SAN quorum.

```
cluster1::> event log show -messagename scsiblade.*
Time                Node             Severity      Event
------------------- ---------------- ------------- ---------------------------
8/13/2013 14:03:51  node0            INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
8/13/2013 14:03:51  node1            INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
```

Example

This example shows a two-node cluster after performing an upgrade. Each node shows a previous out of SAN quorum event message from when the node was upgraded. However, the most recent event message for each node shows that both nodes are presently in SAN quorum.

```
cluster1::> event log show -messagename scsiblade.*
Time                Node             Severity      Event
------------------- ---------------- ------------- ---------------------------
8/13/2013 15:37:51  node1            INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
8/13/2013 15:31:26  node1            CRITICAL      scsiblade.out.of.quorum: This node ...
8/13/2013 15:31:26  node1            INFORMATIONAL scsiblade.ics.trace.dumpfile: The ...
8/13/2013 15:31:26  node1            INFORMATIONAL scsiblade.ics.trace.dumpfile: The ...
8/13/2013 15:30:43  node0            INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
8/13/2013 15:24:16  node0            CRITICAL      scsiblade.out.of.quorum: This node ...
8/13/2013 15:24:16  node0            INFORMATIONAL scsiblade.ics.trace.dumpfile: The ...
8/13/2013 15:24:16  node0            INFORMATIONAL scsiblade.ics.trace.dumpfile: The ...
```

Verifying cluster and SVM health

Before and after you upgrade, revert, or downgrade a cluster, you should verify that the nodes are healthy and eligible to participate in the cluster, and the aggregates and volumes are online.

Steps

1. Verify that the nodes in the cluster are online and are eligible to participate in the cluster:
   `cluster show`

Example

```
cluster1::> cluster show
Node                  Health  Eligibility
--------------------- ------- ------------
node0                 true    true
node1                 true    true
```

If any node is unhealthy or ineligible, check EMS logs for errors and take corrective action. For more information about EMS messages, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

2. Determine if any disk drives are broken, undergoing maintenance, or reconstructing:

To check for...	Do this...
Broken disks	**a.** Display any broken disks: **`storage disk show -state broken`** **b.** Remove or replace any broken disks.
Disks undergoing maintenance or reconstructing	**a.** Display any disks in maintenance, pending, or reconstructing states: **`storage disk show -state maintenance|pending|reconstructing`** **b.** Wait for the maintenance or reconstruction operation to complete before proceeding.

3. To verify that all aggregates are online, display the state of physical and logical storage, including storage aggregates:

`storage aggregate show -state !online`

This command displays the aggregates that are *not* online.

Example

```
cluster1::> storage aggregate show -state !online
There are no entries matching your query.
```

For more information about managing aggregates, see the *Clustered Data ONTAP Physical Storage Management Guide*.

4. To verify that all volumes are online, display any volumes *not* online:

`volume show -state !online`

Example

```
cluster1::> volume show -state !online
There are no entries matching your query.
```

For more information about managing volumes, see the *Clustered Data ONTAP Logical Storage Management Guide*.

Enabling and reverting LIFs to home ports

During a reboot, some LIFs might have been migrated to their assigned failover ports. Before and after you upgrade, revert, or downgrade a cluster, you must enable and revert any LIFs that are not on their home ports.

About this task

The `network interface revert` command reverts a LIF that is not currently on its home port back to its home port, provided that the home port is operational. A LIF's home port is specified when the LIF is created; you can determine the home port for a LIF by using the `network interface show` command.

Steps

1. Display the status of all LIFs:

   `network interface show`

   #### Example

   This example displays the status of all LIFs for a Storage Virtual Machine (SVM, formerly known as Vserver).

   ```
   cluster1::> network interface show -vserver vs0
               Logical    Status     Network            Current       Current Is
   Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
   ----------- ---------- ---------- ------------------ ------------- ------- ----
   vs0
               data001    down/down  192.0.2.120/24     node0         e0e     true
               data002    down/down  192.0.2.121/24     node0         e0f     true
               data003    down/down  192.0.2.122/24     node0         e2a     true
               data004    down/down  192.0.2.123/24     node0         e2b     true
               data005    down/down  192.0.2.124/24     node0         e0e     false
               data006    down/down  192.0.2.125/24     node0         e0f     false
               data007    down/down  192.0.2.126/24     node0         e2a     false
               data008    down/down  192.0.2.127/24     node0         e2b     false
   8 entries were displayed.
   ```

   If any LIFs appear with a `Status Admin` status of `down` or with an `Is home` status of `false`, continue with the next step.

2. Enable the data LIFs:

   `network interface modify {-role data} -status-admin up`

   #### Example

   ```
   cluster1::> network interface modify {-role data} -status-admin up
   8 entries were modified.
   ```

3. Revert LIFs to their home ports:

   `network interface revert *`

Example

This command reverts all LIFs back to their home ports and changes all LIF home statuses to `true`.

```
cluster1::> network interface revert *
8 entries were acted on.
```

4. Verify that all LIFs are in their home ports:
   `network interface show`

Example

This example shows that all LIFs for SVM vs0 are on their home ports.

```
cluster1::> network interface show -vserver vs0
            Logical    Status     Network            Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
----------- ---------- ---------- ------------------ ------------- ------- ----
vs0
            data001      up/up    192.0.2.120/24     node0         e0e     true
            data002      up/up    192.0.2.121/24     node0         e0f     true
            data003      up/up    192.0.2.122/24     node0         e2a     true
            data004      up/up    192.0.2.123/24     node0         e2b     true
            data005      up/up    192.0.2.124/24     node1         e0e     true
            data006      up/up    192.0.2.125/24     node1         e0f     true
            data007      up/up    192.0.2.126/24     node1         e2a     true
            data008      up/up    192.0.2.127/24     node1         e2b     true
8 entries were displayed.
```

Identifying active CIFS sessions that should be terminated

Before performing a minor nondisruptive upgrade or downgrade within the Data ONTAP 8.2 release family, you should identify and gracefully terminate any CIFS sessions that are not continuously available.

About this task

Continuously available CIFS shares, which are accessed by Hyper-V clients using the SMB3 protocol, do not need to be terminated before upgrading or downgrading.

Steps

1. Identify any established CIFS sessions that are not continuously available:
   `vserver cifs session show -continuously-available !Yes -instance`

   This command displays detailed information about any CIFS sessions that have no continuous availability.

Example

```
cluster1::> vserver cifs session show -continuously-available !Yes -
instance
```

```
                        Node: node1
                     Vserver: vs1
                  Session ID: 1
               Connection ID: 4160072788
Incoming Data LIF IP Address: 198.51.100.5
      Workstation IP address: 203.0.113.20
    Authentication Mechanism: NTLMv2
                Windows User: CIFSLAB\user1
                   UNIX User: nobody
                 Open Shares: 1
                  Open Files: 2
                  Open Other: 0
              Connected Time: 8m 39s
                   Idle Time: 7m 45s
            Protocol Version: SMB2_1
      Continuously Available: No
1 entry was displayed.
```

Each of the sessions identified by this command should be terminated before proceeding with the Data ONTAP upgrade or downgrade.

2. If necessary, identify the files that are open for each CIFS session that you identified:
   `vserver cifs session file show -session-id` ***`session_ID`***

Example

```
cluster1::> vserver cifs session file show -session-id 1

Node:       node1
Vserver:    vs1
Connection: 4160072788
Session:    1
File    File      Open Hosting
Continuously
ID      Type      Mode Volume          Share                 Available
------- --------- ---- --------------- ---------------------
------------
1       Regular   rw   vol10           homedirshare          No
Path: \TestDocument.docx
2       Regular   rw   vol10           homedirshare          No
Path: \file1.txt
2 entries were displayed.
```

Verifying the system time

You should verify that NTP is configured, and that the time is synchronized across the cluster.

About this task

For more information about managing the system time, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

Steps

1. Use the `system services ntp server show` command to verify that each node is associated with an NTP server.

   Example

```
cluster1::> system services ntp server show
Node   Server               Version
------ -------------------- -----------
node0
       ntp1.example.com     max
       ntp2.example.com     max
node1
       ntp1.example.com     max
       ntp2.example.com     max
node2
       ntp1.example.com     max
       ntp2.example.com     max
node3
       ntp1.example.com     max
       ntp2.example.com     max
```

2. Verify that each node has the same date and time:

If you are running Data ONTAP...	Enter the following command...
8.1.x	**`system node date show`**
8.2.x	**`cluster date show`**

   Example

```
cluster1::> cluster date show
Node      Date                Timezone
--------- ------------------- -------------------------
node0     4/6/2013 20:54:38   GMT
node1     4/6/2013 20:54:38   GMT
node2     4/6/2013 20:54:38   GMT
node3     4/6/2013 20:54:38   GMT
4 entries were displayed.
```

Checking for back-end configuration errors

Before downgrading gateways to an earlier release, you need to run `storage errors show` to determine whether there are any back-end configuration errors.

Steps

1. Enter the following command:

`storage array config show`

2. Proceed based on the result of Step 1, as follows:

If...	Then...
If the `storage array config show` output *does not* instruct you to run `storage errors show`	Proceed with downgrading.
If the `storage array config show` output *does* instruct you to run `storage errors show`	Continue to the next step.

   You are instructed to run the `storage errors show` command if Data ONTAP detects a back-end configuration error that would prevent Data ONTAP and the back-end storage array from operating together properly.

3. Enter the following command:

   `storage errors show`

   The `storage errors show` command provides details, at the array LUN level, as the following example shows:

   Example

   ```
   IBM_1742_1
   ----------
   NAME (Serial #):  This Array LUN is only available on one path.
   Proper configuration requires two paths.
   ```

4. Fix the problem indicated by `storage errors show`, then downgrade your system.

   The *FlexArray Virtualization Installation Requirements and Reference Guide* contains explanations about errors shown in `storage errors show` output and provides information about how to fix them.

Determining the current software version on each node before downgrading

Ensure that you are running on the minimum software version allowed for the downgrade process by running the `system node image show` command.

Step

1. Determine the current software version:

   `system node image show`

Example

```
cluster1::> system node image show
                 Is      Is                 Install
Node     Image   Default Current Version    Date
-------- ------- ------- ------- --------   -------------------
node0
         image1  true    true    8.2.1      11/27/2013 13:52:22
         image2  false   false   8.2.0      3/25/2013 12:37:36
node1
         image1  true    true    8.2.1      11/27/2013 13:55:22
         image2  false   false   8.2.0      3/25/2013 12:41:16
4 entries were displayed.
```

Obtaining Data ONTAP software images

You must copy a software image from the N series support website (accessed and navigated as described in *Websites*) to an HTTP server on your network so that nodes can access the images by using the `system node image update` command.

About this task

To upgrade, revert, or downgrade the cluster to the target release of Data ONTAP, you need access to software images. Software images, firmware version information, and the latest firmware for your platform model are available on the N series support website (accessed and navigated as described in *Websites* on page 7). Note the following important information:

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

Steps

1. Locate the target Data ONTAP software on the N series support website (accessed and navigated as described in *Websites* on page 7).
2. Copy the software image (for example, `821_q_image.tgz`) from the N series support website (accessed and navigated as described in *Websites* on page 7) to the directory on the HTTP server from which the image will be served.

Installing Data ONTAP software images in a cluster

If needed, install the target Data ONTAP 8.x image software package, but leave the default set to the current Data ONTAP 8.x version.

Before you begin

You must have obtained the Data ONTAP software images.

Steps

1. Choose one of the following options based on your requirements:

If you want to...	Enter this command...
Download, but not install, the software image	**system node image get -node * -package *location* -replace-package true -background true** This command downloads the software image to all of the nodes simultaneously. To download the image to each node one at a time, do not specify the `-background` parameter.
Install a previously downloaded software image	**system node image update -node * -package file:///mroot/etc/software/*image_name* -background true** Note the following considerations for this command: • If you are currently running Data ONTAP 8.2, the `-package` parameter only requires you to enter the image name; you do not need to enter the full directory path. • If you are unsure of the image name to install, you can view a list of previously downloaded software images by using the `system node image package show` command. • This command installs the software image on all of the nodes simultaneously. To install the image on each node one at a time, do not specify the `-background` parameter.
Download and install the software image in the same operation	**system node image update -node * -package *location* -replace-package true -background true** This command downloads and installs the software image on all of the nodes simultaneously. To download and install the image on each node one at a time, do not specify the `-background` parameter.

2. Verify that the software image is downloaded and installed on each node:

   system node image show-update-progress -node *

   This command displays the current status of the software image download and installation. You should continue to run this command until all nodes report a `Run Status` of `Exited`, and an `Exit Status` of `Success`.

   #### Example

   This example shows a 2-node cluster in which the software image has been downloaded and installed successfully on both nodes:

```
cluster1::> system node image show-update-progress -node *
There is no update/install in progress
Status of most recent operation:
        Run Status:     Exited
        Exit Status:    Success
        Phase:          Run Script
        Exit Message:   Installation complete. image2 updated on node node0.
There is no update/install in progress
Status of most recent operation:
        Run Status:     Exited
        Exit Status:    Success
        Phase:          Run Script
        Exit Message:   Installation complete. image2 updated on node node1.
2 entries were acted on.
```

Preparing SnapMirror relationships for a nondisruptive upgrade or downgrade

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

About this task

For more information about SnapMirror operations, see the SnapMirror man pages and the *Clustered Data ONTAP Data Protection Guide*.

Steps

1. Use the `snapmirror show` command to determine the destination path for each SnapMirror relationship.
2. For each destination volume, suspend future SnapMirror transfers:

   `snapmirror quiesce -destination-path` ***`destination`***

   If there are no active transfers for the SnapMirror relationship, this command sets its status to `Quiesced`. If the relationship has active transfers, the status is set to `Quiescing` until the transfer is completed, and then the status becomes `Quiesced`.

   #### Example

   If you are upgrading from Data ONTAP 8.1, this example quiesces transfers involving the destination volume `vol1` from Storage Virtual Machine (SVM) `vs0` and cluster `cluster1`:

   ```
   cluster1::> snapmirror quiesce -destination-path cluster1://vs0/vol1
   ```

   #### Example

   If you are downgrading within the Data ONTAP 8.2 release family, this example quiesces transfers involving the destination volume `vol1` from SVM `vs0`:

   ```
   cluster1::> snapmirror quiesce -destination-path vs0:vol1
   ```

3. Verify that all SnapMirror relationships are quiesced:

   `snapmirror show -status !Quiesced`

This command displays any SnapMirror relationships that are *not* quiesced.

Example

This example shows that all SnapMirror relationships are quiesced:

```
cluster1::> snapmirror show -status !Quiesced
There are no entries matching your query.
```

4. If any SnapMirror relationships are currently transferring, do one of the following options:

Option	Description
Wait for the transfers to complete before performing the Data ONTAP upgrade.	Once each transfer completes, the relationship changes to `Quiesced` status.
Stop the transfers by entering the following command: **`snapmirror abort -destination-path`** ***`destination`*** **`-h`** **Note:** You must use the `-foreground true` parameter if you are aborting load-sharing mirror transfers.	This command stops the SnapMirror transfer and restores the destination volume to the last Snapshot copy that was successfully transferred. The relationship is set to `Quiesced` status.

Ensuring that no jobs are running

You must verify the status of cluster jobs before upgrading or downgrading to a different Data ONTAP release. If any aggregate, volume, NDMP (dump or restore), or Snapshot jobs (such as create, delete, move, modify, replicate, and mount jobs) are running or queued, allow the jobs to finish successfully or stop the queued entries.

Steps

1. Review the list of any running or queued aggregate, volume, or Snapshot jobs:

   `job show`

Example

```
cluster1::> job show
                            Owning
Job ID Name                 Vserver    Node           State
------ -------------------- ---------- -------------- ----------
8629   Vol Reaper           cluster1   -              Queued
       Description: Vol Reaper Job
8630   Certificate Expiry Check
                            cluster1   -              Queued
       Description: Certificate Expiry Check
8632   CLUSTER BACKUP AUTO daily
                            cluster1   -              Queued
       Description: Cluster Backup Job
8633   CLUSTER BACKUP AUTO weekly
                            cluster1   -              Queued
       Description: Cluster Backup Job
```

```
9944   SnapMirrorDaemon_7_2147484678
                            cluster1   node1          Dormant
       Description: Snapmirror Daemon for 7_2147484678
18277  CLUSTER BACKUP AUTO 8hour
                            cluster1   -              Queued
       Description: Cluster Backup Job
18377  SnapMirror Service Job
                            cluster1   node0          Dormant
       Description: SnapMirror Service Job
18379  Network Consistency Diagnostic - weekly
                            cluster1   node0          Queued
       Description: Network Consistency Checker
18385  Network Consistency Diagnostic - weekly
                            cluster1    node1         Queued
       Description: Network Consistency Checker
9 entries were displayed
```

2. Delete any running or queued aggregate, volume, or Snapshot copy jobs:

   `job delete -id` ***`job_id`***

   #### Example

   ```
   cluster1::> job delete -id 8629
   ```

3. Ensure that no aggregate, volume, or Snapshot jobs are running or queued:

   `job show`

   #### Example

   In this example, all running and queued jobs have been deleted.

   ```
   cluster1::> job show
                               Owning
   Job ID Name                 Vserver    Node           State
   ------ -------------------- ---------- -------------- ----------
   9944   SnapMirrorDaemon_7_2147484678
                               cluster1   node1          Dormant
          Description: Snapmirror Daemon for 7_2147484678
   18377  SnapMirror Service Job
                               cluster1   node0          Dormant
          Description: SnapMirror Service Job
   2 entries were displayed
   ```

Performing the Data ONTAP downgrade process

To downgrade a cluster to an earlier Data ONTAP release in the same release family, you must install target images, address downgrade issues, and change the default boot image.

Before you begin

You must complete the downgrade preparation phase before you perform the downgrade procedures.

About this task

You can perform either a nondisruptive downgrade, in which the cluster remains online and continues to serve data during the downgrade, or a disruptive downgrade, in which the cluster is taken offline.

Choices

- *Downgrading a Data ONTAP cluster nondisruptively* on page 149
- *Downgrading a Data ONTAP cluster disruptively* on page 152

Downgrading a Data ONTAP cluster nondisruptively

The nondisruptive downgrade method has several steps: setting the default boot image for each node in the cluster, initiating a failover operation on each node in an HA pair, updating the "failed" node, initiating giveback, and then repeating the process for each HA pair in the cluster.

Before you begin

The cluster must consist of two or more nodes.

Steps

1. Verify that the target Data ONTAP 8.2 software is installed:

`system node image show`

Example

This example shows that version 8.2.0 is installed as the alternate image on both nodes.

```
cluster1::> system node image show
                 Is      Is                 Install
Node     Image   Default Current Version    Date
-------- ------- ------- ------- --------   -------------------
node0
         image1  true    true    8.2.1      11/25/2013 12:37:36
         image2  false   false   8.2.0      4/22/2013 13:52:22
node1
         image1  true    true    8.2.1      11/25/2013 12:41:16
         image2  false   false   8.2.0      4/22/2013 13:55:22
4 entries were displayed.
```

For more information about installing the target Data ONTAP software image, see *Installing Data ONTAP 8.x software images in a cluster*.

2. Change the current default boot image to 8.2.x:

`system node image modify {-iscurrent true} -isdefault false`

This command identifies any functionality in the current release that is not supported in the earlier release. If any of these conditions are found, you must address them according to the instructions provided in the command output before you can proceed.

Example

This example shows that the default boot image will be changed to 8.2.0:

```
cluster1::> system node image modify {-iscurrent true} -isdefault false
2 entries were modified.
```

3. Redisplay the default boot image:

   `system node image show`

   Example

   This example shows version 8.2.0 as the default image on both nodes:

```
cluster1::> system node image show
                 Is      Is                  Install
Node     Image   Default Current Version     Date
-------- ------- ------- ------- --------    -------------------
node0
         image1  false   true    8.2.1       11/25/2013 12:37:36
         image2  true    false   8.2.0       4/22/2013 13:52:22
node1
         image1  false   true    8.2.1       11/25/2013 12:41:16
         image2  true    false   8.2.0       4/22/2013 13:55:22
4 entries were displayed.
```

4. Ensure that storage failover is enabled and possible:

   `storage failover show`

   Example

   This example shows that storage failover is enabled and possible on nodes node0 and node1:

```
cluster1::> storage failover show
                                      Takeover InterConn
Node           Partner        Enabled Possible Up        State
-------------- -------------- ------- -------- --------- -----------
node0          node1          true    true     true      connected
node1          node0          true    true     true      connected
2 entries were displayed.
```

5. If the cluster consists of only two nodes (a single HA pair), ensure that cluster HA is configured:

   `cluster ha show`

   Example

```
cluster1::> cluster ha show
High Availability Configured: true
```

6. Disable automatic giveback on both nodes of the HA pair if it is enabled:

   `storage failover modify -node` ***`nodename`*** **`-auto-giveback false`**

If the cluster is a two-node cluster, a message is displayed warning you that disabling automatic giveback will prevent the management cluster services from going online in the event of a alternating-failure scenario. Enter `y` to continue.

7. Migrate LIFs away from the node that will be taken over during the downgrade:

   `network interface migrate-all -node` ***`nodename`***

   Data LIFs for SAN protocols are not migrated. As long as these LIFs exist on each node in the cluster, data can be served through alternate paths during the upgrade process.

8. Initiate a takeover:

   `storage failover takeover -ofnode` ***`nodename`***

   Do not specify the parameter `-option immediate`, because a normal takeover is required for the node that is being taken over to boot onto the alternate software image.

   The node that is taken over boots up to the `waiting for giveback` state.

   > **Note:** If AutoSupport is enabled, an AutoSupport message is sent indicating that the node is out of cluster quorum. You can safely ignore this notification and proceed with the downgrade.

9. Verify that the takeover was successful:

   `storage failover show`

   ### Example

   This example shows that the takeover was successful. Node node0 is in the `Waiting for giveback` state, and its partner is `In takeover`.

   ```
   cluster1::> storage failover show
                                 Takeover
   Node           Partner        Possible State Description
   -------------- -------------- -------- -------------------------------------
   node0          node1          -        Waiting for giveback (HA mailboxes)
   node1          node0          false    In takeover
   2 entries were displayed.
   ```

10. Wait 8 minutes to ensure the following conditions:

    - Client multipathing (if deployed) is stabilized.
    - Clients are recovered from the pause in I/O that occurs during takeover.
      The recovery time is client-specific and may take longer than 8 minutes depending on the characteristics of the client applications.

11. Return the root aggregate to the partner node:

    `storage failover giveback –ofnode` ***`nodename`***

    > **Attention:** The giveback is not initiated, an error message is returned, and an event is generated if any conditions such as the following are detected:
    >
    > - Long-running operations (such as ASUP generation)
    > - Operations that cannot be restarted (such as aggregate creation)

- Error conditions (such as disk connectivity mismatch between the nodes)

If giveback is not initiated, complete the following steps:

a. Address the "veto" condition described in the error message, ensuring that any identified operations are terminated gracefully.

b. Reenter the giveback command:

```
storage failover giveback -ofnode nodename
```

Alternatively, you can analyze the messages and events for relevance to your environment. If you determine that the veto conditions are not significant, you can override the giveback veto by entering the following command:

```
storage failover giveback –ofnode nodename -override-vetoes true
```

For more information about determining whether you can safely override the veto, see the *Clustered Data ONTAP High-Availability Configuration Guide*.

This first returns the root aggregate to the partner node and then, after that node has finished booting, returns the non-root aggregates.

Before proceeding to the next step, ensure that all of the aggregates have been returned to node B. This process takes approximately 10 minutes.

12. Ensure that all of the aggregates have been returned:

```
storage failover show-giveback
```

After all of the aggregates are returned, the newly booted node begins to serve data to the clients.

13. Repeat Steps *7* on page 151 through *12* on page 152 on the other node in the HA pair.

14. Reenable automatic giveback on both nodes if it was previously enabled:

```
storage failover modify -node nodename -auto-giveback true
```

15. If you are downgrading a cluster that contains multiple HA pairs, ensure that the cluster is in quorum and that services are running before downgrading the next pair of nodes.

Downgrading a Data ONTAP cluster disruptively

If you can take your cluster offline to downgrade Data ONTAP, or if you have a single-node cluster, you can use the disruptive downgrade method. This method has several steps: disabling storage failover for each HA pair, updating the software on each node in the cluster, and then reenabling storage failover.

About this task

During a disruptive downgrade, each node acts as a single-node cluster. Any failures in the node will cause a data outage.

Steps

1. Verify that the target Data ONTAP 8.2 software is installed:

   `system node image show`

   #### Example

   This example shows that version 8.2.0 is installed as the alternate image on both nodes.

   ```
   cluster1::> system node image show
                    Is      Is                 Install
   Node     Image   Default Current Version    Date
   -------- ------- ------- ------- --------   -------------------
   node0
            image1  true    true    8.2.1      11/25/2013 12:37:36
            image2  false   false   8.2.0      4/22/2013 13:52:22
   node1
            image1  true    true    8.2.1      11/25/2013 12:41:16
            image2  false   false   8.2.0      4/22/2013 13:55:22
   4 entries were displayed.
   ```

   For more information about installing the target Data ONTAP software image, see *Installing Data ONTAP 8.x software images in a cluster*.

2. Change the current default boot image to 8.2.x:

   `system node image modify {-iscurrent true} -isdefault false`

   This command identifies any functionality in the current release that is not supported in the earlier release. If any of these conditions are found, you must address them according to the instructions provided in the command output before you can proceed.

   #### Example

   This example shows that the default boot image will be changed to 8.2.0:

   ```
   cluster1::> system node image modify {-iscurrent true} -isdefault false
   2 entries were modified.
   ```

3. Redisplay the default boot image:

   `system node image show`

   #### Example

   This example shows version 8.2.0 as the default image on both nodes:

   ```
   cluster1::> system node image show
                    Is      Is                 Install
   Node     Image   Default Current Version    Date
   -------- ------- ------- ------- --------   -------------------
   node0
            image1  false   true    8.2.1      11/25/2013 12:37:36
            image2  true    false   8.2.0      4/22/2013 13:52:22
   node1
            image1  false   true    8.2.1      11/25/2013 12:41:16
            image2  true    false   8.2.0      4/22/2013 13:55:22
   4 entries were displayed.
   ```

4. Perform one of the following actions:

If the cluster consists of...	Do this...
One node	Continue to the next step.

Two nodes	**a.** Disable cluster high availability:
	`cluster ha modify -configured false`
	b. Disable storage failover for the HA pair:
	`storage failover modify -node * -enabled false`

More than two nodes	Disable storage failover for each HA pair in the cluster: **`storage failover modify -node * -enabled false`**

5. Reboot each node in the cluster:

   `system node reboot -node` ***`nodename`***

   If the cluster consists of more than one node, you can reboot the nodes simultaneously.

   Each node boots the new Data ONTAP image. The Data ONTAP login prompt appears, indicating that the reboot process is complete.

6. When each node has rebooted with the new Data ONTAP image, confirm that the new Data ONTAP 8.2.x software is running:

   `system node image show`

   Example

   This example shows version 8.2.1 as the current version on both nodes:

```
cluster1::> system node image show
                 Is      Is                 Install
Node     Image   Default Current Version    Date
-------- ------- ------- ------- --------   -------------------
node0
         image1  true    true    8.2.1      11/22/2013 13:52:22
         image2  false   false   8.1.2      10/25/2012 12:37:36
node1
         image1  true    true    8.2.1      11/22/2013 13:55:22
         image2  false   false   8.1.2      10/25/2012 12:41:16
4 entries were displayed.
```

7. Enable storage failover for each HA pair in the cluster:

   `storage failover modify -node * -enabled true`

8. If the cluster consists of two nodes, enable cluster high availability:

   `cluster ha modify -configured true`

Completing post-downgrade tasks

After you downgrade the cluster to an earlier version of Data ONTAP 8.x, you should ensure that the cluster is functioning correctly.

Verifying that the cluster is in quorum

Before and after you perform an upgrade, reversion, or downgrade, you must ensure that all nodes are participating in a replicated database (RDB) quorum and that all rings are in the quorum. You must also verify that the per-ring quorum master is the same for all nodes.

About this task

For more information about cluster replication rings and RDB quorums, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

Steps

1. Set the privilege level to advanced:

   `set -privilege advanced`

   Enter **`y`** to continue.

2. Display each RDB process:

To display this RDB process...	Enter this command...
Management application	**`cluster ring show -unitname mgmt`**
Volume location database	**`cluster ring show -unitname vldb`**
Virtual-Interface manager	**`cluster ring show -unitname vifmgr`**
SAN management daemon	**`cluster ring show -unitname bcomd`**

   #### Example

   This example shows the volume location database process for a cluster running Data ONTAP 8.1.x:

   ```
   cluster1::*> cluster ring show -unitname vldb

   Node   UnitName Epoch DB Epoch DB Trnxs Master
   ------ -------- ----- -------- -------- ---------
   node0  vldb     154   154      14847    node0
   node1  vldb     154   154      14847    node0
   node2  vldb     154   154      14847    node0
   node3  vldb     154   154      14847    node0
   4 entries were displayed.
   ```

Example

This example shows the volume location database process for a cluster running Data ONTAP 8.2.x:

```
cluster1::*> cluster ring show -unitname vldb
Node      UnitName Epoch    DB Epoch DB Trnxs Master    Online
--------- -------- -------- -------- -------- --------- ---------
node0     vldb     154      154      14847    node0     master
node1     vldb     154      154      14847    node0     secondary
node2     vldb     154      154      14847    node0     secondary
node3     vldb     154      154      14847    node0     secondary
4 entries were displayed.
```

For each process, verify the following configuration details:

- The relational database epoch and database epochs match for each node.
- The per-ring quorum master is the same for all nodes.
  Note that each ring might have a different quorum master.

3. Return to the admin privilege level:

   `set -privilege admin`

4. If you are operating in a SAN environment, verify that each node is in a SAN quorum:

   `event log show -messagename scsiblade.*`

   The most recent `scsiblade` event message for each node should indicate that the scsi-blade is in quorum. During the upgrade, reversion, or downgrade process, each node will temporarily fall out of SAN quorum. Therefore, if you are verifying the SAN quorum after completing an upgrade, reversion, or downgrade, you may notice critical event messages warning you that the nodes were previously out of SAN quorum.

   If a node is out of SAN quorum, you can use the `storage failover takeover` and `storage failover giveback` commands to perform a planned takeover and giveback with the node's high-availability partner and bring the node back into SAN quorum.

Example

In this example, both nodes in the cluster are in SAN quorum.

```
cluster1::> event log show -messagename scsiblade.*
Time                Node             Severity      Event
------------------- ---------------- ------------- ---------------------------
8/13/2013 14:03:51  node0            INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
8/13/2013 14:03:51  node1            INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
```

Example

This example shows a two-node cluster after performing an upgrade. Each node shows a previous out of SAN quorum event message from when the node was upgraded. However, the most recent event message for each node shows that both nodes are presently in SAN quorum.

```
cluster1::> event log show -messagename scsiblade.*
Time                Node             Severity      Event
------------------- ---------------- ------------- ---------------------------
8/13/2013 15:37:51  node1            INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
8/13/2013 15:31:26  node1            CRITICAL      scsiblade.out.of.quorum: This node ...
8/13/2013 15:31:26  node1            INFORMATIONAL scsiblade.ics.trace.dumpfile: The ...
8/13/2013 15:31:26  node1            INFORMATIONAL scsiblade.ics.trace.dumpfile: The ...
8/13/2013 15:30:43  node0            INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
8/13/2013 15:24:16  node0            CRITICAL      scsiblade.out.of.quorum: This node ...
8/13/2013 15:24:16  node0            INFORMATIONAL scsiblade.ics.trace.dumpfile: The ...
8/13/2013 15:24:16  node0            INFORMATIONAL scsiblade.ics.trace.dumpfile: The ...
```

Verifying cluster and SVM health

Before and after you upgrade, revert, or downgrade a cluster, you should verify that the nodes are healthy and eligible to participate in the cluster, and the aggregates and volumes are online.

Steps

1. Verify that the nodes in the cluster are online and are eligible to participate in the cluster:

   `cluster show`

   #### Example

   ```
   cluster1::> cluster show
   Node                  Health  Eligibility
   --------------------- ------- ------------
   node0                 true    true
   node1                 true    true
   ```

   If any node is unhealthy or ineligible, check EMS logs for errors and take corrective action. For more information about EMS messages, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

2. Determine if any disk drives are broken, undergoing maintenance, or reconstructing:

To check for...	Do this...
Broken disks	**a.** Display any broken disks: **`storage disk show -state broken`** **b.** Remove or replace any broken disks.
Disks undergoing maintenance or reconstructing	**a.** Display any disks in maintenance, pending, or reconstructing states: **`storage disk show -state maintenance|pending|reconstructing`** **b.** Wait for the maintenance or reconstruction operation to complete before proceeding.

3. To verify that all aggregates are online, display the state of physical and logical storage, including storage aggregates:

   `storage aggregate show -state !online`

This command displays the aggregates that are *not* online.

Example

```
cluster1::> storage aggregate show -state !online
There are no entries matching your query.
```

For more information about managing aggregates, see the *Clustered Data ONTAP Physical Storage Management Guide*.

4. To verify that all volumes are online, display any volumes *not* online:

   `volume show -state !online`

   #### Example

   ```
   cluster1::> volume show -state !online
   There are no entries matching your query.
   ```

   For more information about managing volumes, see the *Clustered Data ONTAP Logical Storage Management Guide*.

Enabling and reverting LIFs to home ports

During a reboot, some LIFs might have been migrated to their assigned failover ports. Before and after you upgrade, revert, or downgrade a cluster, you must enable and revert any LIFs that are not on their home ports.

About this task

The `network interface revert` command reverts a LIF that is not currently on its home port back to its home port, provided that the home port is operational. A LIF's home port is specified when the LIF is created; you can determine the home port for a LIF by using the `network interface show` command.

Steps

1. Display the status of all LIFs:

   `network interface show`

   #### Example

   This example displays the status of all LIFs for a Storage Virtual Machine (SVM, formerly known as Vserver).

   ```
   cluster1::> network interface show -vserver vs0
               Logical    Status     Network            Current       Current Is
   Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
   ----------- ---------- ---------- ------------------ ------------- ------- ----
   vs0
   ```

```
            data001    down/down  192.0.2.120/24     node0         e0e     true
            data002    down/down  192.0.2.121/24     node0         e0f     true
            data003    down/down  192.0.2.122/24     node0         e2a     true
            data004    down/down  192.0.2.123/24     node0         e2b     true
            data005    down/down  192.0.2.124/24     node0         e0e     false
            data006    down/down  192.0.2.125/24     node0         e0f     false
            data007    down/down  192.0.2.126/24     node0         e2a     false
            data008    down/down  192.0.2.127/24     node0         e2b     false
8 entries were displayed.
```

If any LIFs appear with a `Status Admin` status of `down` or with an `Is home` status of `false`, continue with the next step.

2. Enable the data LIFs:

   `network interface modify {-role data} -status-admin up`

   #### Example

   ```
   cluster1::> network interface modify {-role data} -status-admin up
   8 entries were modified.
   ```

3. Revert LIFs to their home ports:

   `network interface revert *`

   #### Example

   This command reverts all LIFs back to their home ports and changes all LIF home statuses to `true`.

   ```
   cluster1::> network interface revert *
   8 entries were acted on.
   ```

4. Verify that all LIFs are in their home ports:

   `network interface show`

   #### Example

   This example shows that all LIFs for SVM vs0 are on their home ports.

   ```
   cluster1::> network interface show -vserver vs0
               Logical    Status     Network            Current       Current Is
   Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
   ----------- ---------- ---------- ------------------ ------------- ------- ----
   vs0
               data001      up/up    192.0.2.120/24     node0         e0e     true
               data002      up/up    192.0.2.121/24     node0         e0f     true
               data003      up/up    192.0.2.122/24     node0         e2a     true
               data004      up/up    192.0.2.123/24     node0         e2b     true
               data005      up/up    192.0.2.124/24     node1         e0e     true
               data006      up/up    192.0.2.125/24     node1         e0f     true
               data007      up/up    192.0.2.126/24     node1         e2a     true
               data008      up/up    192.0.2.127/24     node1         e2b     true
   8 entries were displayed.
   ```

Verifying client access (CIFS and NFS)

For the configured protocols, test access from CIFS and NFS clients to verify that the cluster is accessible.

Resuming SnapMirror operations

After completing a nondisruptive upgrade or downgrade, you must resume any SnapMirror relationships that were suspended.

Before you begin

Existing SnapMirror relationships must have been suspended by using the `snapmirror quiesce` command, and the cluster must have been nondisruptively upgraded or downgraded.

Steps

1. Resume transfers for each SnapMirror relationship that was previously quiesced:

   `snapmirror resume *`

2. Verify that the SnapMirror operations have resumed:

   `snapmirror show`

   #### Example

```
cluster1::> snapmirror show

Source            Destination  Mirror  Relationship  Total             Last
Path        Type  Path         State   Status        Progress  Healthy Updated
----------- ---- ------------ ------- -------------- --------- ------- --------
cluster1-vs1:dp_src1
            DP   cluster1-vs2:dp_dst1
                              Snapmirrored
                                      Idle           -         true    -
cluster1-vs1:xdp_src1
            XDP  cluster1-vs2:xdp_dst1
                              Snapmirrored
                                      Idle           -         true    -
cluster1://cluster1-vs1/ls_src1
            LS   cluster1://cluster1-vs1/ls_mr1
                              Snapmirrored
                                      Idle           -         true    -
                 cluster1://cluster1-vs1/ls_mr2
                              Snapmirrored
                                      Idle           -         true    -
4 entries were displayed.
```

   For each SnapMirror relationship, verify that the Relationship Status is "Idle". If the status is "Transferring", wait for the SnapMirror transfer to complete, and then reenter the command to verify that the status has changed to "Idle".

After you finish

For each SnapMirror relationship that is configured to run on a schedule, you should verify that the first scheduled SnapMirror transfer completes successfully.

Considerations for downgrading the SP firmware

If your current SP firmware version is not supported for the Data ONTAP release you are downgrading or reverting to, you must install a supported SP firmware version for the earlier Data ONTAP release.

After the Data ONTAP revert or downgrade process is complete, you must take the action to install an SP firmware version that is supported for the Data ONTAP version you reverted or downgraded to. For information about downloading and installing an SP firmware version, see the N series support website (accessed and navigated as described in *Websites* on page 7).

Reinstalling the required gateway license after revert or downgrade

The licensing scheme for gateways changed in Data ONTAP 8.2. Data ONTAP 8.2 and later requires the V_StorageAttach license package for gateways to be able to access LUNs on storage arrays. Depending on the type of reversion you do, you might need to install the gateway license key for a prior release.

A license key is required for each gateway being reverted or downgraded.

If...	Then the gateway license requirement is...
The gateway was upgraded to Data ONTAP 8.2 from the 8.1.x release family and you want to revert to a release in the 8.1.x release family again	You do not have to install the gateway license for the release to which the system is being reverted. Data ONTAP remembers the gateway license from when it was upgraded from 8.1.x to 8.2; Data ONTAP reinstalls those licenses when Data ONTAP is downgraded to 8.1.x.
Data ONTAP 8.2 is the first release of Data ONTAP installed on the system and you want to revert to a release in the 8.1.x release family	You must manually install the license key for gateways that is supported for the release to which the system is being reverted. Contact your sales representative to obtain the appropriate license key for your system and release.
You want to downgrade the system from Data ONTAP 8.2.x to an earlier release in the 8.2 release family	You do not have to install the V_StorageAttach license package; Data ONTAP remembers the appropriate license for the release family.

For information about how to install licenses, see the *Clustered Data ONTAP System Administration Guide for SVM Administrators*.

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 a client perspective.

How upgrades impact service availability

You can review the factors that can affect the availability of cluster services before you begin the upgrade.

The following factors impact 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 upgrades when downtime is scheduled, but it does not affect services in nondisruptive upgrades.
- 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 cluster services. Make sure that you review upgrade options carefully to determine the best method of upgrading for maintaining optimal service availability.

Related concepts

Identifying potential upgrade issues on page 24

Updating firmware on page 84

Considerations for services and protocols during upgrades

In general, services based on stateless protocols—such as NFS, FC, and iSCSI—are less susceptible to service interruptions during upgrades than session-oriented protocols—such as CIFS and NDMP.

During an upgrade, each node in the cluster must be rebooted (by initiating an HA configuration takeover and giveback) to load the new software. Services based on stateless protocols usually remain available during the nondisruptive upgrade.

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 cluster, if the client's timeout period is greater than the disruption period on the cluster (for example, the amount of time a reboot or HA configuration giveback takes), the client does not perceive a disruption of cluster services.

In session-oriented protocols, there is no concept of timeout to protect the service from disruption. If session-oriented cluster 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 NAS and SAN protocols generally do not experience adverse effects during upgrades if the clients are configured according to recommended guidelines.

If you are using stateless protocols, consider the following:

- NFS hard mounts

  No adverse behavior is experienced on the clients during upgrade. Clients might receive some messages similar to the following until the node 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 an NFS operation reaches a timeout using soft mounts.

  Situations that can cause frequent timeouts include nondisruptive upgrades or any takeover or giveback event in an HA configuration.

  In general, soft mounts should be used only when reading solely from a disk; even then, understand that any soft mount is unreliable.
- SAN protocols

  No adverse behavior is experienced on FC or iSCSI clients if they are configured according to recommended guidelines.

  For compatibility and configuration information about FC and iSCSI products, see the N series Interoperability Matrices website (accessed and navigated as described in *Websites* on page 7).

Considerations for session-oriented protocols

Clusters and session-oriented protocols might cause adverse effects on clients and applications in certain areas during upgrades.

If you are using session-oriented protocols, consider the following:

- CIFS

Hyper-V over SMB supports nondisruptive operations (NDO). If you configured a Hyper-V over SMB solution, Hyper-V and the contained virtual machines remain online and provide continuous availability during the Data ONTAP upgrade.

For all other CIFS configurations, client sessions are terminated. You should direct users to end their sessions before you upgrade.

- NDMP

  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 node reboots and there is new disk firmware present, the affected drives are automatically and sequentially taken offline, and the node 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 node 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 nodeshell `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 cluster operation.

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

- Degraded aggregates are on the node.
- Disks needing a firmware update are present in an aggregate or plex that is in an offline state.

Automatic background disk firmware updates resume when these conditions are addressed. For more information about determining aggregate status and state, see the *Clustered Data ONTAP Physical Storage Management Guide*.

Copyright and trademark information

Copyright ©1994 - 2014 NetApp, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 2014 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 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, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Mars, 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, Snap Creator, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, VelocityStak, 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

NA 210-06392_A0, Printed in USA

SC27-6284-01