

Cluster administration ONTAP 9

NetApp August 25, 2025

This PDF was generated from https://docs.netapp.com/usen/ontap/concept_administration_overview.html on August 25, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Cluster administration.	1
Cluster management with System Manager	1
Learn about cluster administration with ONTAP System Manager	1
Use System Manager to access an ONTAP cluster.	2
Configure protocols on your ONTAP cluster	3
Enable new features by adding license keys with ONTAP System Manager	4
Download a cluster configuration with ONTAP System Manager	4
Assign tags to a cluster with ONTAP System Manager	5
View and submit support cases with ONTAP System Manager	6
Manage the maximum capacity limit of a storage VM in ONTAP System Manager	6
Monitor cluster, tier, and SVM capacity in ONTAP System Manager.	8
View hardware configurations to determine problems with ONTAP System Manager	10
Manage nodes using ONTAP System Manager	16
License management.	18
ONTAP licensing overview	18
Download NetApp license files (NLF) from NetApp Support Site	21
Install NetApp licenses in ONTAP	22
Manage ONTAP licenses	23
License types and licensed method	25
Commands for managing licenses in ONTAP	26
Cluster management with the CLI	27
Learn about cluster administration with the ONTAP CLI	27
Cluster and SVM administrators	27
Access the cluster by using the CLI (cluster administrators only)	
Use the ONTAP command-line interface	40
Record an ONTAP CLI session and manage the recorded sessions	54
Cluster management (cluster administrators only)	55
Manage nodes	60
Configure the SP/BMC network	83
Manage nodes remotely using the SP/BMC	
Manage ON IAP cluster time (cluster administrators only).	
Manage the banner and MOTD	
Manage ONTAP jobs and job schedules	129
Back up and restore cluster conligurations (cluster administrators only).	1/1
Disk and tier management	1/2
	143
Disks and ONTAP local tiers	143
Manage local tiers	195
Manage RAID configurations	100 016
Manage NAD configurations	∠10 ววว
FabricDool tier management	ՀՀՀ
	201 027
	201

Requirements for using ONTAP FabricPool.	237
Tier data efficiently with ONTAP FabricPool policies	242
Learn about ONTAP FabricPool configuration and management tasks	246
Configure FabricPool	246
Manage FabricPool	264
Manage FabricPool mirrors	286
ONTAP commands for managing FabricPool resources	293
SVM data mobility	296
SVM data mobility overview	296
Migrate an SVM	303
Monitor volume migration	305
Pause and resume SVM migration	305
Cancel an SVM migration	306
Manually cut over clients	306
Manually remove source SVM after client cutover.	307
HA pair management	307
Learn about HA pair management in ONTAP clusters	307
Learn about hardware-assisted takeovers in ONTAP clusters	308
Learn about automatic takeover and giveback in ONTAP clusters	309
ONTAP automatic takeover commands	313
ONTAP automatic giveback commands.	313
ONTAP manual takeover commands	316
ONTAP manual giveback commands	318
Testing takeover and giveback in ONTAP clusters	321
ONTAP commands for monitoring an HA pair	323
ONTAP commands for enabling and disabling storage failover.	327
Halt or reboot ONTAP nodes without initiating takeover in two-node clusters	327
REST API management with System Manager	330
REST API management with System Manager	331
Accessing the REST API log	331

Cluster administration

Cluster management with System Manager

Learn about cluster administration with ONTAP System Manager

System Manager is an HTML5-based graphical management interface that enables you to use a web browser to manage storage systems and storage objects (such as disks, volumes, and storage tiers) and perform common management tasks related to storage systems.

- System Manager is included with ONTAP software as a web service, enabled by default, and accessible by using a browser.
- (\mathbf{i})
- The name of System Manager has changed beginning with ONTAP 9.6. In ONTAP 9.5 and earlier it was called OnCommand System Manager. Beginning with ONTAP 9.6 and later, it is called System Manager.
- If you are using the classic System Manager (available only in ONTAP 9.7 and earlier), refer to System Manager Classic (ONTAP 9.0 to 9.7)

Using the System Manager Dashboard, you can view at-a-glance information about important alerts and notifications, the efficiency and capacity of storage tiers and volumes, the nodes that are available in a cluster, the status of the nodes in an HA pair, the most active applications and objects, and the performance metrics of a cluster or a node.

With System Manager in ONTAP 9.7 and later releases, you can perform many common tasks such as the following:

- Create a cluster, configure a network, and set up support details for the cluster.
- Configure and manage storage objects, such as disks, local tiers, volumes, qtrees, and quotas.
- · Configure protocols, such as SMB and NFS, and provision file sharing.
- Configure protocols such as FC, FCoE, NVMe, and iSCSI for block access.
- Create and configure network components, such as subnets, broadcast domains, data and management interfaces, and interface groups.
- Set up and manage mirroring and vaulting relationships.
- Perform cluster management, storage node management, and storage virtual machine (storage VM) management operations.
- Create and configure storage VMs, manage storage objects associated with storage VMs, and manage storage VM services.
- Monitor and manage high-availability (HA) configurations in a cluster.
- Configure service processors to remotely log in, manage, monitor, and administer the node, regardless of the state of the node.

System Manager terminology

System Manager uses different terminology than the CLI for some ONTAP key functionality.

- Local tier: A set of physical solid-state drives or hard-disk drives you store your data on. You might know these as aggregates. In fact, if you use the ONTAP CLI, you will still see the term *aggregate* used to represent a local tier.
- **Cloud tier**: Storage in the cloud used by ONTAP when you want to have some of your data off premises for one of several reasons. If you are thinking of the cloud part of a FabricPool, you've already figured it out. And if you are using a StorageGRID system, your cloud might not be off premises at all. (A cloud-like experience on premises is called a *private cloud*.)
- **Storage VM**: A virtual machine running within ONTAP that provides storage and data services to your clients. You might know this as an *SVM* or a *vserver*.
- **Network interface**: Address and properties assigned to a physical network port. You might know this as a *logical interface (LIF)*.
- **Pause**: An action that halts operations. Before ONTAP 9.8, you might have referred to *quiesce* in other versions of System Manager.

Use System Manager to access an ONTAP cluster

If you prefer to use a graphic interface instead of the command-line interface (CLI) for accessing and managing a cluster, you can do so by using System Manager, which is included with ONTAP as a web service, is enabled by default, and is accessible by using a browser.

Beginning with ONTAP 9.12.1, System Manager is fully integrated with BlueXP.



With BlueXP, you can manage your hybrid multicloud infrastructure from a single control plane while retaining the familiar System Manager dashboard.

See System Manager integration with BlueXP.

About this task

You can use a cluster management network interface (LIF) or node management network interface (LIF) to access System Manager. For uninterrupted access to System Manager, you should use a cluster management network interface (LIF).

Before you begin

- You must have a cluster user account that is configured with the "admin" role and the "http" and "console" application types.
- You must have enabled cookies and site data in the browser.

Steps

- 1. Point the web browser to the IP address of the cluster management network interface:
 - ° If you are using IPv4: https://cluster-mgmt-LIF
 - o If you are using IPv6: https://[cluster-mgmt-LIF]



Only HTTPS is supported for browser access of System Manager.

If the cluster uses a self-signed digital certificate, the browser might display a warning indicating that the certificate is not trusted. You can either acknowledge the risk to continue the access or install a Certificate Authority (CA) signed digital certificate on the cluster for server authentication.

2. **Optional:** If you have configured an access banner by using the CLI, then read the message that is displayed in the **Warning** dialog box, and choose the required option to proceed.

This option is not supported on systems on which Security Assertion Markup Language (SAML) authentication is enabled.

- If you do not want to continue, click Cancel, and close the browser.
- If you want to continue, click **OK** to navigate to the System Manager login page.
- 3. Log in to System Manager by using your cluster administrator credentials.



Beginning with ONTAP 9.11.1, when you log in to System Manager, you can specify the locale. The locale specifies certain localization settings, such as language, currency, time and date format, and similar settings. For ONTAP 9.10.1 and earlier, the locale for System Manager is detected from the browser. To change the locale for System Manager, you have to change the locale of the browser.

- 4. **Optional**: Beginning with ONTAP 9.12.1, you can specify your preference for the appearance of System Manager:
 - a. In the upper right corner of System Manager, click **I** to manage user options.
 - b. Position the System Theme toggle switch to your preference:

Toggle position	Appearance setting
📧 (left)	Light theme (Light background with dark text)
OS (center)	Default to the theme preference that was set for the operating system's applications (usually the theme setting for the browser that is used to access System Manager).
(right)	Dark theme (Dark background with light text)

Related information

Managing access to web services

Accessing a node's log, core dump, and MIB files by using a web browser

Configure protocols on your ONTAP cluster

Depending on the licenses enabled on your cluster, you can enable the desired protocols on your cluster. You then create network interfaces using which you can access the storage.

About this task

This procedure applies to FAS, AFF, and ASA systems. If you have an ASA r2 system (ASA A1K, ASA A90, ASA A70, ASA A50, ASA A30, ASA A20, or ASA C30), follow these steps to use System Manger to set up an ONTAP cluster. ASA r2 systems provide a simplified ONTAP experience specific to SAN-only customers.

- 1. Select **Dashboard** and then click **Configure Protocols**.
 - $\,\circ\,$ To enable NAS protocols, select NFS or SMB.
 - $\,\circ\,$ To enable SAN protocols, select **iSCSI** or **FC**.
 - To enable NVMe protocols, select **NVMe**.

Enable new features by adding license keys with ONTAP System Manager

In releases earlier than ONTAP 9.10.1, ONTAP features are enabled with license keys, and features in ONTAP 9.10.1 and later are enabled with a NetApp license file. You can add license keys and NetApp license files using System Manager.

Beginning with ONTAP 9.10.1, you use System Manager to install a NetApp License File to enable multiple licensed features all at once. Using a NetApp License File simplifies license installation because you no longer have to add separate feature license keys. You download the NetApp License File from the NetApp Support Site.

If you already have license keys for some features and you are upgrading to ONTAP 9.10.1, you can continue to use those license keys.

Steps

- 1. Select Cluster > Settings.
- 2. Under Licenses, select \rightarrow .
- 3. Select **Browse**. Choose the NetApp License File you downloaded.
- 4. If you have license keys you want to add, select Use 28-character license keys and enter the keys.

Download a cluster configuration with ONTAP System Manager

Beginning with ONTAP 9.11.1, you can use System Manager to download some configuration details about the cluster and its nodes. This information can be used for inventory management, hardware replacement, and lifecycle activities. This information is especially useful to sites that do not send AutoSupport (ASUP) data.

Cluster configuration details include the cluster name, cluster ONTAP version, cluster management LIF, volume, and LIF counts.

Node configuration details include the node name, system serial number, system ID, system model, ONTAP version, MetroCluster information, SP/BMC network information, and encryption configuration information.

Steps

- 1. Click Cluster > Overview.
- ^{2.} Click **More** to display the drop-down menu.
- 3. Select Download configuration.
- 4. Select the HA pairs, then click Download.

The configuration is downloaded as an Excel spreadsheet.

• The first sheet contains cluster details.

• The other sheets contain node details.

Assign tags to a cluster with ONTAP System Manager

Beginning with ONTAP 9.14.1, you can use System Manager to assign tags to a cluster to identify objects as belonging to a category, such as projects or cost centers.

About this task

You can assign a tag to a cluster. First, you need to define and add the tag. Then, you can also edit or delete the tag.

Tags can be added when you create a cluster, or they can be added later.

You define a tag by specifying a key and associating a value to it using the format "key:value". For example: "dept:engineering" or "location:san-jose".

The following should be considered when you create tags:

- Keys have a minimum length of one character and cannot be null. Values can be null.
- A key can be paired with multiple values by separating the values with a comma, for example, "location:san-jose,toronto"
- Tags can be used for multiple resources.
- Keys must start with a lowercase letter.

Steps

To manage tags, performing the following steps:

1. In System Manager, click **Cluster** to view the overview page.

The tags are listed in the Tags section.

2. Click Manage Tags to modify existing tags or add new ones.

You can add, edit, or delete the tags.

To perform this action	Perform these steps
Add a tag	a. Click Add Tag.b. Specify a key and its value or values (separate multiple values with commas).c. Click Save.
Edit a tag	a. Modify the content in the Key and Values (optional) fields.b. Click Save.
Delete a tag	a. Click 🧃 next to the tag you want to delete.

View and submit support cases with ONTAP System Manager

Beginning with ONTAP 9.9.1, you can view support cases from Active IQ Digital Advisor (also known as Digital Advisor) associated with the cluster. You can also copy cluster details that you need to submit a new support case on the NetApp Support Site. Beginning with ONTAP 9.10.1, you can enable telemetry logging, which helps support personnel troubleshoot problems.



To receive alerts about firmware updates, you must be registered with Active IQ Unified Manager. Refer to Active IQ Unified Manager documentation resources.

Steps

1. In System Manager, select Support.

A list of open support cases associated with this cluster is displayed.

- 2. Click on the following links to perform procedures:
 - Case Number: See details about the case.
 - **Go to NetApp Support Site**: Navigate to the **My AutoSupport** page on the NetApp Support Site to view knowledge base articles or submit a new support case.
 - View My Cases: Navigate to the My Cases page on the NetApp Support Site.
 - View Cluster Details: View and copy information you will need when you submit a new case.

Enable telemetry logging

Beginning with ONTAP 9.10.1, you can use System Manager to enable telemetry logging. When telemetry logging is allowed, messages that are logged by System Manager are given a specific telemetry identifier that indicates the exact process that triggered the message. All messages that are issued relating to that process have the same identifier, which consists of the name of the operational workflow and a number (for example "add-volume-1941290").

If you experience performance problems, you can enable telemetry logging, which allows support personnel to more easily identify the specific process for which a message was issued. When telemetry identifiers are added to the messages, the log file is only slightly enlarged.

Steps

- 1. In System Manager, select **Cluster > Settings**.
- 2. In UI Settings section, click the check box for Allow telemetry logging.

Manage the maximum capacity limit of a storage VM in ONTAP System Manager

Beginning with ONTAP 9.13.1, you can use System Manager to enable a maximum capacity limit for a storage VM and set a threshold to trigger alerts when the used storage reaches a certain percentage of the maximum capacity.

Enable a maximum capacity limit for a storage VM

Beginning with ONTAP 9.13.1, you can specify the maximum capacity that can be allocated for all volumes in a storage VM. You can enable the maximum capacity when you add a storage VM or when you edit an existing

storage VM.

Steps

- 1. Select Storage > Storage VMs.
- 2. Perform one of the following:
 - To add a storage VM, click + Add
 - To edit a storage VM, click inext to the name of the storage VM, and then click Edit.
- 3. Enter or modify the settings for the storage VM, and select the check box labeled "Enable maximum capacity limit".
- 4. Specify the maximum capacity size.
- 5. Specify the percentage of the maximum capacity you want to use as a threshold to trigger alerts.
- 6. Click Save.

Edit the maximum capacity limit of a storage VM

Beginning with ONTAP 9.13.1, you can edit the maximum capacity limit of an existing storage VM, if the maximum capacity limit has been enabled already.

Steps

- 1. Select Storage > Storage VMs.
- 2. Click inext to the name of the storage VM, and then click Edit.

The check box labeled "Enable maximum capacity limit" is already checked.

3. Perform one of the following steps:

Action	Steps
Disable the maximum capacity limit	 Uncheck the check box. Click Save.
Modify the maximum capacity limit	 Specify the new maximum capacity size. (You cannot specify a size that is less than the already allocated space in the storage VM.)
	Specify the new percentage of the maximum capacity you want to use as a threshold to trigger alerts.
	3. Click Save .

Related information

- View the maximum capacity limit of a storage VM
- · Capacity measurements in System Manager
- Manage SVM capacity limits

Monitor cluster, tier, and SVM capacity in ONTAP System Manager

Using System Manager, you can monitor how much storage capacity has been used and how much is still available for a cluster, a local tier, or a storage VM.

With each version of ONTAP, System Manager provides more robust capacity monitoring information:

- Beginning with ONTAP 9.13.1, you can enable a maximum capacity limit for a storage VM and set a threshold to trigger alerts when the used storage reaches a certain percentage of the maximum capacity.
- Beginning with ONTAP 9.12.1, System Manager displays the amount of committed capacity for a local tier.
- Beginning with ONTAP 9.10.1, System Manager lets you view historical data about the cluster's capacity and projections about how much capacity will be used or available in the future. You can also monitor the capacity of local tiers and volumes.



Measurements of used capacity are displayed differently depending on your ONTAP version. Learn more in Capacity measurements in System Manager.

View the capacity of a cluster

You can view capacity measurements for a cluster on the Dashboard in System Manager.

Before you begin

To view data related to the capacity in the cloud, you must have an account with Digital Advisor and be connected.

Steps

- 1. In System Manager, click Dashboard.
- 2. In the Capacity section, you can view the following:
 - Total used capacity of the cluster
 - · Total available capacity of the cluster
 - · Percentages of used and available capacity.
 - Ratio of data reduction.
 - Amount of capacity used in the cloud.
 - History of capacity usage.
 - Projection of capacity usage



In System Manager, capacity representations do not account for root storage tier (aggregate) capacities.

3. Click the chart to view more details about the capacity of the cluster.

Capacity measurements are shown in two bar charts:

- The top chart displays the physical capacity: the size of physical used, reserved, and available space.
- The bottom chart displays the logical capacity: the size of client data, snapshots, and clones, and the total logical used space.

Below the bar charts are measurements for data reduction:

- Data reduction ratio for only the client data (snapshots and clones are not included).
- Overall data reduction ratio.

For more information, see Capacity measurements in System Manager.

View the capacity of a local tier

You can view details about the capacity of local tiers. Beginning with ONTAP 9.12.1, the **Capacity** view also includes the amount of committed capacity for a local tier, enabling you to determine whether you need to add capacity to the local tier to accommodate the committed capacity and avoid running out of free space.

Steps

- 1. Click Storage > Tiers.
- 2. Select the name of the local tier.
- 3. On the **Overview** page, in the **Capacity** section, the capacity is show in a bar chart with three measurements:
 - Used and reserved capacity
 - Available capacity
 - Committed capacity (beginning with ONTAP 9.12.1)
- 4. Click the chart to view details about the capacity of the local tier.

Capacity measurements are shown in two bar charts:

- The top bar chart displays physical capacity: the size of physical used, reserved, and available space.
- The bottom bar chart displays logical capacity: the size of client data, snapshots, and clones, and the total of logical used space.

Below the bar charts are measurements ratios for data reduction:

- Data reduction ratio for only the client data (snapshots and clones are not included).
- · Overall data reduction ratio.

For more information, see Capacity measurements in System Manager.

Optional actions

- If the committed capacity is larger than the capacity of the local tier, you might consider adding capacity to the local tier before it runs out of free space. See Add capacity to a local tier (add disks to an aggregate).
- You can also view the storage that specific volumes use in the local tier by selecting the Volumes tab.

View the capacity of the volumes in a storage VM

You can view how much storage is used by the volumes in a storage VM and how much capacity is still available. The total measurement of used and available storage is called "capacity across volumes".

Steps

1. Select Storage > Storage VMs.

- 2. Click on the name of the storage VM.
- 3. Scroll to the **Capacity** section, which shows a bar chart with the following measurements:
 - Physical used: Sum of physical used storage across all volumes in this storage VM.
 - Available: Sum of available capacity across all volumes in this storage VM.
 - Logical used: Sum of logical used storage across all volumes in this storage VM.

For more details about the measurements, see Capacity measurements in System Manager.

View the maximum capacity limit of a storage VM

Beginning with ONTAP 9.13.1, you can view the maximum capacity limit of a storage VM.

Before you begin

You must enable the maximum capacity limit of a storage VM before you can view it.

Steps

1. Select Storage > Storage VMs.

You can view the maximum capacity measurements in two ways:

- In the row for the storage VM, view the **Maximum Capacity** column which contains a bar chart that shows the used capacity, available capacity, and maximum capacity.
- Click the name of the storage VM. On the **Overview** tab, scroll to view the maximum capacity, allocated capacity, and capacity alert threshold values in the left column.

Related information

- Edit the maximum capacity limit of a storage VM
- · Capacity measurements in System Manager

View hardware configurations to determine problems with ONTAP System Manager

Beginning with ONTAP 9.8, you can use System Manager to view the configuration of hardware on your network and determine the health of your hardware systems and cabling configurations.

Steps

To view hardware configurations, perform the following steps:

- 1. In System Manager, select Cluster > Hardware.
- 2. Hover your mouse over components to view status and other details.

You can view various types of information:

- Information about controllers
- Information about disk shelves
- Information about storage switches
- 3. Beginning with ONTAP 9.12.1, you can view cabling information in System Manager. Click the Show

Cables check box to view cabling, then hover over a cable to view its connectivity information.

• Information about cabling

Information about controllers

You can view the following:

Nodes

- You can view the front and rear views.
- For models with an internal disk shelf, you can also view the disk layout in the front view.
- You can view the following platforms:

Platfor	Support	Supported in System Manager in ONTAP version…							
m	9.16.1	9.15.1	9.14.1	9.13.1	9.12.1	9.11.1	9.10.1	9.9.1	9.8 (previe w mode only)
AFF A20	Yes								
AFF A30	Yes								
AFF A50	Yes								
AFF A70		Yes							
AFF A90		Yes							
AFF A1K		Yes							
AFF A150		Yes	Yes	Yes					
AFF A220		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
AFF A250		Yes	Yes	Yes	Yes	Yes	Yes	Yes	
AFF A300		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
AFF A320		Yes	Yes	Yes	Yes	Yes	Yes	Yes	
AFF A400		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

AFF A700		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
AFF A700s		Yes	Yes	Yes	Yes	Yes	Yes	Yes	
AFF A800		Yes	Yes	Yes	Yes	Yes	Yes	Yes	
AFF C30	Yes								
AFF C60	Yes								
AFF C80	Yes								
AFF C190		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
AFF C250		Yes	Yes	Yes	Yes *	Yes *	Yes *		
AFF C400		Yes	Yes	Yes	Yes *	Yes *	Yes *		
AFF C800		Yes	Yes	Yes	Yes *	Yes *	Yes *		
ASA A150		Yes	Yes	Yes					
ASA A250		Yes	Yes	Yes					
ASA A400		Yes	Yes	Yes					
ASA A800		Yes	Yes	Yes					
ASA A900		Yes	Yes	Yes					
ASA C250		Yes	Yes	Yes					

Yes	Yes	Yes					
Yes	Yes	Yes					
Yes							
Yes							
Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Yes	Yes	Yes	Yes	Yes			
Yes	Yes	Yes	Yes	Yes			
Yes	Yes	Yes	Yes	Yes			
Yes	Yes	Yes	Yes	Yes			
Yes	Yes	Yes	Yes	Yes			
Yes	Yes	Yes	Yes	Yes			
	Yes Yes Yes Yes Yes Yes Yes Yes	YesYesYes	YesYesYesYesIIYesIIYes	YesYesYesYesIIYesIIYes	YesYesYesIYesIIIYesIIIYes	YesYesYesImage: Constraint of the sector of	YesYesYesImage: Constraint of the sector of

Ports

- You will see a port highlighted in red if it is down.
- When you hover over the port, you can view the status of a port and other details.
- You cannot view console ports.

Notes:

- Beginning with ONTAP 9.11.1, you will see SAS ports highlighted in red only if they are in an error state or if a cabled port that is being used goes offline. The ports appear in white if they are offline and uncabled.
- For ONTAP 9.10.1 and earlier, you will see SAS ports highlighted in red when they are disabled.

FRUs

Information about FRUs appears only when the state of a FRU is non-optimal.

- · Failed PSUs in nodes or chassis.
- High temperatures detected in nodes.
- Failed fans on the nodes or chassis.

Adapter cards

- Cards with defined part number fields display in the slots if external cards have been inserted.
- Ports display on the cards.
- For a supported card, you can view images of that card. If the card is not in the list of supported part numbers, then a generic graphic appears.

Information about disk shelves

You can view the following:

Disk shelves

- You can display the front and rear views.
- You can view the following disk shelf models:

If your system is running	Then you can use System Manager to view
ONTAP 9.9.1 and later	All shelves that have <i>not</i> been designated as "end of service" or "end of availability"
ONTAP 9.8	DS4243, DS4486, DS212C, DS2246, DS224C, and NS224

Shelf ports

- You can view port status.
- You can view remote port information if the port is connected.

Shelf FRUs

• PSU failure information displays.

Information about storage switches

You can view the following:

Storage switches

- The display shows switches that act as storage switches used to connect shelves to nodes.
- Beginning with ONTAP 9.9.1, System Manager displays information about a switch that acts as both a storage switch and a cluster, which can also be shared between nodes of an HA pair.
- The following information displays:
 - Switch name
 - IP address
 - Serial number
 - SNMP version
 - · System version
- You can view the following storage switch models:

If your system is running	Then you can use System Manager to view…
ONTAP 9.11.1 or later	Cisco Nexus 3232C Cisco Nexus 9336C-FX2 NVIDIA SN2100
ONTAP 9.10.1 and 9.9.1	Cisco Nexus 3232C Cisco Nexus 9336C-FX2
ONTAP 9.8	Cisco Nexus 3232C

Storage switch ports

- The following information displays:
 - Identity name
 - Identity index
 - State
 - Remote connection
 - Other details

Information about cabling

Beginning with ONTAP 9.12.1, you can view the following cabling information:

- · Cabling between controllers, switches, and shelves when no storage bridges are used
- · Connectivity that shows the IDs and MAC addresses of the ports on either end of the cable

Manage nodes using ONTAP System Manager

Using System Manager, you can add nodes to a cluster and rename them. You can also reboot, take over, and give back nodes.

Add nodes to a cluster

You can increase the size and capabilities of your cluster by adding new nodes.

Before you Start

You should have already cabled the new nodes to the cluster.

About this task

There are separate processes for working with System Manager in ONTAP 9.8 and later, or ONTAP 9.7.

ONTAP 9.8 and later procedure

Adding nodes to a cluster with System Manager (ONTAP 9.8 and later)

Steps

1. Select Cluster > Overview.

The new controllers are shown as nodes connected to the cluster network but are not in the cluster.

- 2. Select Add.
 - The nodes are added into the cluster.
 - · Storage is allocated implicitly.

ONTAP 9.7 procedure

Adding nodes to a cluster with System Manager (ONTAP 9.7)

Steps

- 1. Select (Return to classic version).
- 2. Select Configurations > Cluster Expansion.

System Manager automatically discovers the new nodes.

- 3. Select Switch to the new experience.
- 4. Select **Cluster > Overview** to view the new nodes.

Shut down, reboot or edit service processor

When you reboot or shutdown a node, its HA partner automatically executes a takeover.



This procedure applies to FAS, AFF, and ASA systems. If you have an ASA r2 system (ASA A1K, ASA A90, ASA A70, ASA A50, ASA A30, ASA A20, or ASA C30), follow these steps to shutdown and reboot a node. ASA r2 systems provide a simplified ONTAP experience specific to SAN-only customers.

Steps

- 1. Select Cluster > Overview.
- 2. Under Nodes, select .
- 3. Select the node and then select Shut down, Reboot, or Edit Service Processor.

If a node has been rebooted and is waiting for giveback, the **Giveback** option is also available.

If you select **Edit Service Processor**, you can choose **Manual** to input the IP address, subnet mask and gateway, or you can choose **DHCP** for dynamic host configuration.

Rename nodes

Beginning with ONTAP 9.14.1, you can rename a node from the cluster overview page.



This procedure applies to FAS, AFF, and ASA systems. If you have an ASA r2 system (ASA A1K, ASA A90, ASA A70, ASA A50, ASA A30, ASA A20, or ASA C30), follow these steps to rename a node. ASA r2 systems provide a simplified ONTAP experience specific to SAN-only customers.

Steps

- 1. Select Cluster. The cluster overview page displays.
- 2. Scroll down to the Nodes section.
- 3. Next to the node that you want to rename, select , and select **Rename**.
- 4. Modify the node name, and then select **Rename**.

License management

ONTAP licensing overview

A license is a record of one or more software entitlements. Beginning with ONTAP 9.10.1, all licenses are delivered as a NetApp license file (NLF), which is a single file that enables multiple features. Beginning in May 2023, all AFF systems (both A-series and C-series) and FAS systems are sold with either the ONTAP One software suite or the ONTAP Base software suite, and beginning in June 2023, all ASA systems are sold with ONTAP One for SAN. Each software suite is delivered as a single NLF, replacing the separate NLF bundles first introduced in ONTAP 9.10.1.

Licenses included with ONTAP One

ONTAP One contains all available licensed functionality. It contains a combination of the contents of the former Core bundle, Data Protection bundle, Security and Compliance bundle, Hybrid Cloud bundle, and Encryption bundle, as shown in the table. Encryption is not available in restricted countries.

Former bundle name	ONTAP keys included
Core bundle	FlexClone
	SnapRestore
	NFS, SMB, S3
	FC, iSCSI
	NVME-oF

Security and Compliance bundle	Autonomous Ransomware Protection
	МТКМ
	SnapLock
Data Protection bundle	SnapMirror (asynchronous, synchronous, active sync)
	SnapCenter
	SnapMirror S3 for NetApp targets
Hybrid Cloud bundle	SnapMirror cloud
	SnapMirror S3 for non-NetApp targets
Encryption bundle	NetApp Volume Encryption
	Trusted Platform module

Licenses not included with ONTAP One

ONTAP One does not include any of NetApp's cloud-delivered services, including the following:

- BlueXP tiering (previously known as cloud tiering)
- Data Infrastructure Insights
- BlueXP backup
- Data governance

ONTAP One for existing systems

If you have existing systems that are currently under NetApp support but have not been upgraded to ONTAP One, the existing licenses on those systems are still valid and continue to work as expected. For example, if the SnapMirror license is already installed on existing systems, it is not necessary to upgrade to ONTAP One to get a new SnapMirror license. However, if you do not have a SnapMirror license installed on an existing system, the only way to get that license is to upgrade to ONTAP One for an additional fee.

Beginning in June 2023, ONTAP systems using 28-character license keys can also upgrade to the ONTAP One or ONTAP Base compatibility bundle.

Licenses included with ONTAP Base

ONTAP Base is an optional software suite that's an alternative to ONTAP One for ONTAP systems. It is for specific use cases where data protection technologies such as SnapMirror and SnapCenter, as well as security features like Autonomous Ransomware, are not required, such as non-production systems for dedicated test or development environments. Additional licenses cannot be added to ONTAP Base. If you want additional licenses, such as SnapMirror, you must upgrade to ONTAP One.

Former bundle name ONTAP keys included	
--	--

Core bundle	FlexClone
	SnapRestore
	NFS, SMB, S3
	FC, iSCSI
	NVME-oF
Encryption bundle	NetApp Volume Encryption
	Trusted Platform module

Licenses included with ONTAP One for SAN

ONTAP One for SAN is available for ASA A-series and C-series systems. This is the only software suite available for SAN. ONTAP One for SAN contains the following licenses:

ONTAP keys included		
FlexClone		
SnapRestore		
FC, iSCSI		
NVME-oF		
MTKM		
SnapLock		
SnapMirror (asynchronous, synchronous, active sync)		
SnapCenter		
SnapMirror cloud		
NetApp Volume Encryption		
Trusted Platform module		

Other license delivery methods

In ONTAP 8.2 through ONTAP 9.9.1, license keys are delivered as 28-character strings, and there is one key per ONTAP feature. You use the ONTAP CLI to install license keys if you are using ONTAP 8.2 through ONTAP 9.9.1.



ONTAP 9.10.1 supports installing 28-character license keys using System Manager or the CLI. However, if an NLF license is installed for a feature, you cannot install a 28-character license key over the NetApp license file for the same feature. For information about installing NLFs or license keys using System Manager, see Install ONTAP licenses.

Related information

How to get an ONTAP One license when the system has NLFs already

How to verify ONTAP Software Entitlements and related License Keys using the Support Site

NetApp: ONTAP Entitlement Risk Status

Download NetApp license files (NLF) from NetApp Support Site

If your system is running ONTAP 9.10.1 or later, you can upgrade the bundle license files on existing systems by downloading the NLF for ONTAP One or ONTAP Core from the NetApp Support Site.



The SnapMirror cloud and SnapMirror S3 licenses are not included with ONTAP One. They are part of the ONTAP One Compatibility bundle, which you can get for free if you have ONTAP One and request separately.

Steps

You can download ONTAP One license files for systems with existing NetApp license file bundles and for systems with 28-character license keys that have been converted to NetApp license files on systems running ONTAP 9.10.1 and later. For a fee, you can also upgrade systems from ONTAP Base to ONTAP One.

Upgrade existing NLF

1. Contact your NetApp sales team and request the license file bundle you want to upgrade or convert (for example, ONTAP Base to ONTAP One, or Core Bundle and Data Protection bundle to ONTAP One).

When your request is processed, you will receive an email from netappsw@netapp.com with the subject "NetApp Software Licensing Notification for SO# [SO Number]" and the email will include a PDF attachment that includes your license serial number.

- 2. Log in to the NetApp Support Site.
- 3. Select Systems > Software Licenses.
- 4. From the menu, choose **Serial Number**, enter the serial number you received, and click **New Search**.
- 5. Locate the license bundle you want to convert.
- 6. Click **Get NetApp License File** for each license bundle and download the NLFs when they're available.
- 7. Install the ONTAP One file.

Upgrade NLF converted from license key

- 1. Log in to the NetApp Support Site.
- 2. Select Systems > Software Licenses.
- 3. From the menu, choose Serial Number, enter the system serial number, and click New Search.
- 4. Locate the license you want to convert, and in the **Eligibility** column click **Check**.
- 5. In the Check Eligibility form, click Generate Licenses for 9.10.x and later.
- 6. Close the Check Eligibility form.

You will need to wait at least 2 hours for the licenses to generate.

- 7. Repeat Steps 1 through 3.
- 8. Locate the ONTAP One license, click Get NetApp License File, and choose the delivery method.
- 9. Install the ONTAP One file.

Install NetApp licenses in ONTAP

You can install NetApp license files (NLFs) and license keys using System Manager, which is the preferred method for installing NLFs, or you can use the ONTAP CLI to install license keys. In ONTAP 9.10.1 and later, features are enabled with a NetApp license file, and in releases earlier than ONTAP 9.10.1, ONTAP features are enabled with license keys.

Steps

If you have already downloaded NetApp license files or license keys, you can use System Manager or the ONTAP CLI to install NLFs and 28-character license keys.

System Manager - ONTAP 9.8 and later

- 1. Select Cluster > Settings.
- 2. Under Licenses, select \rightarrow .
- 3. Select Browse. Choose the NetApp License File you downloaded.
- 4. If you have license keys you want to add, select Use 28-character license keys and enter the keys.

System Manager - ONTAP 9.7 and earlier

- 1. Select Configuration > Cluster > Licenses.
- 2. Under Licenses, select \rightarrow .
- 3. In the **Packages** window, click **Add**.
- 4. In the **Add License Packages** dialog box, click **Choose Files** to select the NetApp License File that you downloaded, and then click **Add** to upload the file to the cluster.

CLI

1. Add one or more license key:

system license add

The following example installs licenses from the local node "/mroot/etc/lic_file" if the file exists at this location:

cluster1::> system license add -use-license-file true

Learn more about the system license add command in the ONTAP command reference.

Manage ONTAP licenses

You can use System Manager or the ONTAP CLI to view and manage licenses installed on your system, including viewing the license serial number, checking the status of a license, and removing a license.

View details about a license

Steps

How you view details about a license depends on what version of ONTAP you are using and whether you use System Manager or the ONTAP CLI.

System Manager - ONTAP 9.8 and later

- 1. To view details about a specific feature license, select **Cluster > Settings**.
- 2. Under Licenses, select ->.
- 3. Select Features.
- 4. Locate the licensed feature you want to view and select v to view the license details.

System Manager - ONTAP 9.7 and earlier

- 1. Select Configuration > Cluster > Licenses.
- 2. In the Licenses window, perform the appropriate action:
- 3. Click the **Details** tab.

CLI

1. Display details about an installed license:

system license show

Delete a license

System Manager - ONTAP 9.8 and later

- 1. To delete a license, select **Cluster > Settings**.
- 2. Under Licenses, select ->.
- 3. Select Features.
- 4. Select the licensed feature you want to delete and Delete legacy key.

System Manager - ONTAP 9.7 and earlier

- 1. Select Configuration > Cluster > Licenses.
- 2. In the Licenses window, perform the appropriate action:

If you want to	Do this
Delete a specific license package on a node or a master license	Click the Details tab.
Delete a specific license package across all of the nodes in the cluster	Click the Packages tab.

3. Select the software license package that you want to delete, and then click **Delete**.

You can delete only one license package at a time.

4. Select the confirmation check box, and then click Delete.

CLI

1. Delete a license:

system license delete

The following example deletes a license named CIFS and serial number 1-81-00000000000000000000123456 from the cluster:

cluster1::> system license delete -serial-number 1-81-00000000000000000123456 -package CIFS

The following example deletes from the cluster all of the licenses under the installed-license Core Bundle for serial number 123456789:

```
cluster1::> system license delete { -serial-number 123456789
-installed-license "Core Bundle" }
```

Related information

ONTAP CLI commands for managing licenses

License types and licensed method

Understanding license types and the licensed method helps you manage the licenses in a cluster.

License types

A package can have one or more of the following license types installed in the cluster. The system license show command displays the installed license type or types for a package.

• Standard license (license)

A standard license is a node-locked license. It is issued for a node with a specific system serial number (also known as a *controller serial number*). A standard license is valid only for the node that has the matching serial number.

Installing a standard, node-locked license entitles a node to the licensed functionality. For the cluster to use licensed functionality, at least one node must be licensed for the functionality. It might be out of compliance to use licensed functionality on a node that does not have an entitlement for the functionality.

• Site license (site)

A site license is not tied to a specific system serial number. When you install a site license, all nodes in the cluster are entitled to the licensed functionality. The system license show command displays site licenses under the cluster serial number.

If your cluster has a site license and you remove a node from the cluster, the node does not carry the site license with it, and it is no longer entitled to the licensed functionality. If you add a node to a cluster that has a site license, the node is automatically entitled to the functionality granted by the site license.

• Evaluation license (demo)

An evaluation license is a temporary license that expires after a certain period of time (indicated by the system license show command). It enables you to try certain software functionality without purchasing an entitlement. It is a cluster-wide license, and it is not tied to a specific serial number of a node.

If your cluster has an evaluation license for a package and you remove a node from the cluster, the node does not carry the evaluation license with it.

Licensed method

It is possible to install both a cluster-wide license (the site or demo type) and a node-locked license (the license type) for a package. Therefore, an installed package can have multiple license types in the cluster. However, to the cluster, there is only one *licensed method* for a package. The licensed method field of the system license status show command displays the entitlement that is being used for a package. The command determines the licensed method as follows:

- If a package has only one license type installed in the cluster, the installed license type is the licensed method.
- If a package does not have any licenses installed in the cluster, the licensed method is none.

• If a package has multiple license types installed in the cluster, the licensed method is determined in the following priority order of the license type--site, license, and demo.

For example:

- If you have a site license, a standard license, and an evaluation license for a package, the licensed method for the package in the cluster is site.
- If you have a standard license and an evaluation license for a package, the licensed method for the package in the cluster is license.
- If you have only an evaluation license for a package, the licensed method for the package in the cluster is demo.

Commands for managing licenses in ONTAP

You can use the ONTAP CLI system license commands to manage feature licenses for the cluster. You use the system feature-usage commands to monitor feature usage.

Learn more about the commands described in this topic in the ONTAP command reference.

The following table lists some of the common CLI commands for managing licenses and links to the command man pages for additional information.

If you want to	Use this command
Display all packages that require licenses and their current license status, including the following:	system license show-status
• The package name	
 The licensed method 	
 The expiration date, if applicable 	
Display or remove expired or unused licenses	system license clean-up
Display summary of feature usage in the cluster on a per-node basis	system feature-usage show-summary
Display feature usage status in the cluster on a per-node and per-week basis	system feature-usage show-history
Display the status of license entitlement risk for each license package	system license entitlement-risk show

Related information

- ONTAP command reference
- Knowledge Base article: ONTAP 9.10.1 and later licensing overview
- Use System Manager to install a NetApp license file

Cluster management with the CLI

Learn about cluster administration with the ONTAP CLI

You can administer ONTAP systems with the command-line interface (CLI). You can use the ONTAP management interfaces, access the cluster, manage nodes, and much more.

You should use these procedures under the following circumstances:

- You want to understand the range of ONTAP administrator capabilities.
- You want to use the CLI, not System Manager or an automated scripting tool.

Related information

For details about CLI syntax and usage, see the ONTAP command reference documentation.

Cluster and SVM administrators

Learn about ONTAP cluster and SVM administrator roles

Cluster administrators administer the entire cluster and the storage virtual machines (SVMs, formerly known as Vservers) it contains. SVM administrators administer only their own data SVMs.

Cluster administrators can administer the entire cluster and its resources. They can also set up data SVMs and delegate SVM administration to SVM administrators. The specific capabilities that cluster administrators have depend on their access-control roles. By default, a cluster administrator with the "admin" account name or role name has all capabilities for managing the cluster and SVMs.

SVM administrators can administer only their own SVM storage and network resources, such as volumes, protocols, LIFs, and services. The specific capabilities that SVM administrators have depend on the access-control roles that are assigned by cluster administrators.



The ONTAP command-line interface (CLI) continues to use the term *Vserver* in the output, and vserver as a command or parameter name has not changed.

Enable or disable web browser access to ONTAP System Manager

You can enable or disable a web browser's access to System Manager. You can also view the System Manager log.

You can control a web browser's access to System Manager by using vserver services web modify -name sysmgr -vserver <cluster_name> -enabled [true|false].

System Manager logging is recorded in the /mroot/etc/log/mlog/sysmgr.log files of the node that

hosts the cluster management LIF at the time System Manager is accessed. You can view the log files by using a browser. The System Manager log is also included in AutoSupport messages.

Learn about the ONTAP cluster management server

The cluster management server, also called an *admin*SVM, is a specialized storage virtual machine (SVM) implementation that presents the cluster as a single manageable entity. In addition to serving as the highest-level administrative domain, the cluster management server owns resources that do not logically belong with a data SVM.

The cluster management server is always available on the cluster. You can access the cluster management server through the console or cluster management LIF.

Upon failure of its home network port, the cluster management LIF automatically fails over to another node in the cluster. Depending on the connectivity characteristics of the management protocol you are using, you might or might not notice the failover. If you are using a connectionless protocol (for example, SNMP) or have a limited connection (for example, HTTP), you are not likely to notice the failover. However, if you are using a long-term connection (for example, SSH), then you will have to reconnect to the cluster management server after the failover.

When you create a cluster, all of the characteristics of the cluster management LIF are configured, including its IP address, netmask, gateway, and port.

Unlike a data SVM or node SVM, a cluster management server does not have a root volume or host user volumes (though it can host system volumes). Furthermore, a cluster management server can only have LIFs of the cluster management type.

If you run the vserver show command, the cluster management server appears in the output listing for that command.

Types of SVMs in an ONTAP cluster

A cluster consists of four types of SVMs, which help in managing the cluster and its resources and data access to the clients and applications.

A cluster contains the following types of SVMs:

Admin SVM

The cluster setup process automatically creates the admin SVM for the cluster. The admin SVM represents the cluster.

Node SVM

A node SVM is created when the node joins the cluster, and the node SVM represents the individual nodes of the cluster.

• System SVM (advanced)

A system SVM is automatically created for cluster-level communications in an IPspace.

Data SVM

A data SVM represents the data serving SVMs. After the cluster setup, a cluster administrator must create

data SVMs and add volumes to these SVMs to facilitate data access from the cluster.

A cluster must have at least one data SVM to serve data to its clients.



Unless otherwise specified, the term SVM refers to a data (data-serving) SVM.

In the CLI, SVMs are displayed as Vservers.

Access the cluster by using the CLI (cluster administrators only)

Access an ONTAP cluster using the node serial port

You can access the cluster directly from a console that is attached to a node's serial port.

Steps

1. At the console, press Enter.

The system responds with the login prompt.

2. At the login prompt, do one of the following:

To access the cluster with	Enter the following account name
The default cluster account	admin
An alternative administrative user account	username

The system responds with the password prompt.

3. Enter the password for the admin or administrative user account, and then press Enter.

Access an ONTAP cluster using SSH requests

You can issue SSH requests to an ONTAP cluster to perform administrative tasks. SSH is enabled by default.

Before you begin

• You must have a user account that is configured to use ssh as an access method.

The -application parameter of the security login commands specifies the access method for a user account. Learn more about security login in the ONTAP command reference.

• If you use an Active Directory (AD) domain user account to access the cluster, an authentication tunnel for the cluster must have been set up through a CIFS-enabled storage VM, and your AD domain user account must also have been added to the cluster with ssh as an access method and domain as the authentication method.

About this task

- You must use an OpenSSH 5.7 or later client.
- Only the SSH v2 protocol is supported; SSH v1 is not supported.

• ONTAP supports a maximum of 64 concurrent SSH sessions per node.

If the cluster management LIF resides on the node, it shares this limit with the node management LIF.

If the rate of incoming connections is higher than 10 per second, the service is temporarily disabled for 60 seconds.

• ONTAP supports only the AES and 3DES encryption algorithms (also known as *ciphers*) for SSH.

AES is supported with 128, 192, and 256 bits in key length. 3DES is 56 bits in key length as in the original DES, but it is repeated three times.

- When FIPS mode is on, SSH clients should negotiate with Elliptic Curve Digital Signature Algorithm (ECDSA) public key algorithms for the connection to be successful.
- If you want to access the ONTAP CLI from a Windows host, you can use a third-party utility such as PuTTY.
- If you use a Windows AD user name to log in to ONTAP, you should use the same uppercase or lowercase letters that were used when the AD user name and domain name were created in ONTAP.

AD user names and domain names are not case-sensitive. However, ONTAP user names are casesensitive. Case mismatch between the user name created in ONTAP and the user name created in AD results in a login failure.

SSH Authentication options

• Beginning with ONTAP 9.3, you can enable SSH multifactor authentication for local administrator accounts.

When SSH multifactor authentication is enabled, users are authenticated by using a public key and a password.

- Beginning with ONTAP 9.4, you can enable SSH multifactor authentication for LDAP and NIS remote users.
- Beginning with ONTAP 9.13.1, you can optionally add certificate validation to the SSH authentication
 process to enhance login security. To do this, associate an X.509 certificate with the public key that an
 account uses. If you log in using SSH with both an SSH public key and an X.509 certificate, ONTAP checks
 the validity of the X.509 certificate before authenticating with the SSH public key. SSH login is refused if
 that certificate is expired or revoked, and the SSH public key is automatically disabled.
- Beginning with ONTAP 9.14.1, ONTAP administrators can add Cisco Duo two-factor authentication to the SSH authentication process to enhance login security. Upon first login after you enable Cisco Duo authentication, users will need to enroll a device to serve as an authenticator for SSH sessions.
- Beginning with ONTAP 9.15.1, administrators can Configure dynamic authorization to provide additional adaptive authentication to SSH users based on the user's trust score.

Steps

- 1. From a host with access to the ONTAP cluster's network, enter the ssh command in one of the following formats:
 - ° ssh username@hostname_or_IP [command]
 - ° ssh -1 username hostname_or_IP [command]

If you are using an AD domain user account, you must specify <code>username</code> in the format of <code>domainname \ AD_accountname</code> (with double backslashes after the domain name) or

"*domainname*\AD_accountname" (enclosed in double quotation marks and with a single backslash after the domain name).

hostname_or_IP is the host name or the IP address of the cluster management LIF or a node management LIF. Using the cluster management LIF is recommended. You can use an IPv4 or IPv6 address.

command is not required for SSH-interactive sessions.

Examples of SSH requests

The following examples show how the user account named "joe" can issue an SSH request to access a cluster whose cluster management LIF is 10.72.137.28:

```
$ ssh joe@10.72.137.28
Password:
cluster1::> cluster show
Node Health Eligibility
-----
node1 true true
node2 true true
2 entries were displayed.
```

```
$ ssh -l joe 10.72.137.28 cluster show
Password:
Node Health Eligibility
-----
node1 true true
node2 true true
2 entries were displayed.
```

The following examples show how the user account named "john" from the domain named "DOMAIN1" can issue an SSH request to access a cluster whose cluster management LIF is 10.72.137.28:

```
$ ssh DOMAIN1\\john@10.72.137.28
Password:
cluster1::> cluster show
Node Health Eligibility
-----
node1 true true
node2 true true
2 entries were displayed.
```

```
$ ssh -l "DOMAIN1\john" 10.72.137.28 cluster show
Password:
Node Health Eligibility
node1 true true
node2 true true
2 entries were displayed.
```

The following example shows how the user account named "joe" can issue an SSH MFA request to access a cluster whose cluster management LIF is 10.72.137.32:

```
$ ssh joe@10.72.137.32
Authenticated with partial success.
Password:
cluster1::> cluster show
Node Health Eligibility
-----
node1 true true
node2 true true
2 entries were displayed.
```

Related information

Administrator authentication and RBAC

ONTAP SSH login security

Beginning with ONTAP 9.5, you can view information about previous logins, unsuccessful attempts to log in, and changes to your privileges since your last successful login.

Security-related information is displayed when you successfully log in as an SSH admin user. You are alerted about the following conditions:

- The last time your account name was logged in.
- The number of unsuccessful login attempts since the last successful login.
- Whether the role has changed since the last login (for example, if the admin account's role changed from "admin" to "backup.")
- Whether the add, modify, or delete capabilities of the role were modified since the last login.



If any of the information displayed is suspicious, you should immediately contact your security department.

To obtain this information when you login, the following prerequisites must be met:

- · Your SSH user account must be provisioned in ONTAP.
- Your SSH security login must be created.

• Your login attempt must be successful.

Restrictions and other considerations for SSH login security

The following restrictions and considerations apply to SSH login security information:

- The information is available only for SSH-based logins.
- For group-based admin accounts, such as LDAP/NIS and AD accounts, users can view the SSH login information if the group of which they are a member is provisioned as an admin account in ONTAP.

However, alerts about changes to the role of the user account cannot be displayed for these users. Also, users belonging to an AD group that has been provisioned as an admin account in ONTAP cannot view the count of unsuccessful login attempts that occurred since the last time they logged in.

- The information maintained for a user is deleted when the user account is deleted from ONTAP.
- The information is not displayed for connections to applications other than SSH.

Examples of SSH login security information

The following examples demonstrate the type of information displayed after you login.

• This message is displayed after each successful login:

Last Login : 7/19/2018 06:11:32

• These messages are displayed if there have been unsuccessful attempts to login since the last successful login:

```
Last Login : 4/12/2018 08:21:26
Unsuccessful login attempts since last login - 5
```

• These messages are displayed if there have been unsuccessful attempts to login and your privileges were modified since the last successful login:

Last Login : 8/22/2018 20:08:21 Unsuccessful login attempts since last login - 3 Your privileges have changed since last login

Enable Telnet or RSH access to an ONTAP cluster

As a security best practice, Telnet and RSH are disabled by default. To enable the cluster to accept Telnet or RSH requests, you must enable the service in the default management service policy.

Telnet and RSH are not secure protocols; you should consider using SSH to access the cluster. SSH provides a secure remote shell and interactive network session. For more information, refer to Access the cluster using SSH.
About this task

• ONTAP supports a maximum of 50 concurrent Telnet or RSH sessions per node.

If the cluster management LIF resides on the node, it shares this limit with the node management LIF.

If the rate of incoming connections is higher than 10 per second, the service is temporarily disabled for 60 seconds.

• RSH commands require advanced privileges.

ONTAP 9.10.1 or later

Steps

1. Confirm that the RSH or Telnet security protocol is enabled:

security protocol show

- a. If the RSH or Telnet security protocol is enabled, continue to the next step.
- b. If the RSH or Telnet security protocol is not enabled, use the following command to enable it:

security protocol modify -application <rsh/telnet> -enabled true

Learn more about security protocol show and security protocol modify in the ONTAP command reference.

2. Confirm that the management-rsh-server or management-telnet-server service exists on the management LIFs:

network interface show -services management-rsh-server

or

network interface show -services management-telnet-server

Learn more about network interface show in the ONTAP command reference.

- a. If the management-rsh-server or management-telnet-server service exists, continue to the next step.
- b. If the management-rsh-server or management-telnet-server service does not exist, use the following command to add it:

network interface service-policy add-service -vserver cluster1 -policy
default-management -service management-rsh-server

network interface service-policy add-service -vserver cluster1 -policy
default-management -service management-telnet-server

Learn more about network interface service-policy add-service in the ONTAP command reference.

ONTAP 9.9 or earlier

About this task

ONTAP prevents you from changing predefined firewall policies, but you can create a new policy by cloning the predefined mgmt management firewall policy, and then enabling Telnet or RSH under the new policy.

Steps

1. Enter the advanced privilege mode:

set advanced

2. Enable a security protocol (RSH or Telnet):

security protocol modify -application security protocol -enabled true

3. Create a new management firewall policy based on the mgmt management firewall policy:

system services firewall policy clone -policy mgmt -destination-policy
policy-name

4. Enable Telnet or RSH in the new management firewall policy:

system services firewall policy create -policy policy-name -service
security protocol -action allow -ip-list ip address/netmask

To allow all IP addresses, you should specify -ip-list 0.0.0/0

5. Associate the new policy with the cluster management LIF:

```
network interface modify -vserver cluster_management_LIF -lif cluster_mgmt
-firewall-policy policy-name
```

Learn more about network interface modify in the ONTAP command reference.

Access an ONTAP cluster using Telnet requests

You can issue Telnet requests to the cluster to perform administrative tasks. Telnet is disabled by default.

Telnet and RSH are not secure protocols; you should consider using SSH to access the cluster. SSH provides a secure remote shell and interactive network session. For more information, refer to Access the cluster using SSH.

Before you begin

The following conditions must be met before you can use Telnet to access the cluster:

• You must have a cluster local user account that is configured to use Telnet as an access method.

The -application parameter of the security login commands specifies the access method for a user account. Learn more about security login in the ONTAP command reference.

About this task

• ONTAP supports a maximum of 50 concurrent Telnet sessions per node.

If the cluster management LIF resides on the node, it shares this limit with the node management LIF.

If the rate of in-coming connections is higher than 10 per second, the service is temporarily disabled for 60 seconds.

- If you want to access the ONTAP CLI from a Windows host, you can use a third-party utility such as PuTTY.
- RSH commands require advanced privileges.

ONTAP 9.10.1 or later

Steps

1. Confirm that the Telnet security protocol is enabled:

security protocol show

- a. If the Telnet security protocol is enabled, continue to the next step.
- b. If the Telnet security protocol is not enabled, use the following command to enable it:

security protocol modify -application telnet -enabled true

Learn more about security protocol show and security protocol modify in the ONTAP command reference.

2. Confirm that the management-telnet-server service exists on the management LIFs:

network interface show -services management-telnet-server

Learn more about network interface show in the ONTAP command reference.

- a. If the management-telnet-server service exists, continue to the next step.
- b. If the management-telnet-server service does not exist, use the following command to add it:

network interface service-policy add-service -vserver cluster1 -policy
default-management -service management-telnet-server

Learn more about network interface service-policy add-service in the ONTAP command reference.

ONTAP 9.9 or earlier

Before you begin

The following conditions must be met before you can use Telnet to access the cluster:

• Telnet must already be enabled in the management firewall policy that is used by the cluster or node management LIFs so that Telnet requests can go through the firewall.

By default, Telnet is disabled. The system services firewall policy show command with the -service telnet parameter displays whether Telnet has been enabled in a firewall policy. Learn more about system services firewall policy in the ONTAP command reference.

• If you use IPv6 connections, IPv6 must already be configured and enabled on the cluster, and firewall policies must already be configured with IPv6 addresses.

The network options ipv6 show command displays whether IPv6 is enabled. Learn more about network options ipv6 show in the ONTAP command reference. The system services firewall policy show command displays firewall policies.

Steps

1. From an administration host, enter the following command:

telnet hostname_or_IP

hostname_or_IP is the host name or the IP address of the cluster management LIF or a node management LIF. Using the cluster management LIF is recommended. You can use an IPv4 or IPv6 address.

Example of a Telnet request

The following example shows how the user named "joe", who has been set up with Telnet access, can issue a Telnet request to access a cluster whose cluster management LIF is 10.72.137.28:

admin_host\$ telnet 10.72.137.28
Data ONTAP
login: joe
Password:
cluster1::>

Access an ONTAP cluster using RSH requests

You can issue RSH requests to the cluster to perform administrative tasks. RSH is not a secure protocol and is disabled by default.

Telnet and RSH are not secure protocols; you should consider using SSH to access the cluster. SSH provides a secure remote shell and interactive network session. For more information, refer to Access the cluster using SSH.

Before you begin

The following conditions must be met before you can use RSH to access the cluster:

• You must have a cluster local user account that is configured to use RSH as an access method.

The -application parameter of the security login commands specifies the access method for a user account. Learn more about security login in the ONTAP command reference.

About this task

• ONTAP supports a maximum of 50 concurrent RSH sessions per node.

If the cluster management LIF resides on the node, it shares this limit with the node management LIF.

If the rate of incoming connections is higher than 10 per second, the service is temporarily disabled for 60 seconds.

• RSH commands require advanced privileges.

ONTAP 9.10.1 or later

Steps

1. Confirm that the RSH security protocol is enabled:

security protocol show

- a. If the RSH security protocol is enabled, continue to the next step.
- b. If the RSH security protocol is not enabled, use the following command to enable it:

security protocol modify -application rsh -enabled true

Learn more about security protocol show and security protocol modify in the ONTAP command reference.

2. Confirm that the management-rsh-server service exists on the management LIFs:

network interface show -services management-rsh-server

Learn more about network interface show in the ONTAP command reference.

a. If the management-rsh-server service exists, continue to the next step.

b. If the management-rsh-server service does not exist, use the following command to add it:

network interface service-policy add-service -vserver cluster1 -policy
default-management -service management-rsh-server

Learn more about network interface service-policy add-service in the ONTAP command reference.

ONTAP 9.9 or earlier

Before you begin

The following conditions must be met before you can use RSH to access the cluster:

• RSH must already be enabled in the management firewall policy that is used by the cluster or node management LIFs so that RSH requests can go through the firewall.

By default, RSH is disabled. The system services firewall policy show command with the -service rsh parameter displays whether RSH has been enabled in a firewall policy. Learn more about system services firewall policy in the ONTAP command reference.

• If you use IPv6 connections, IPv6 must already be configured and enabled on the cluster, and firewall policies must already be configured with IPv6 addresses.

The network options ipv6 show command displays whether IPv6 is enabled. Learn more about network options ipv6 show in the ONTAP command reference. The system services firewall policy show command displays firewall policies.

Steps

1. From an administration host, enter the following command:

rsh hostname_or_IP -l username:passwordcommand

hostname_or_IP is the host name or the IP address of the cluster management LIF or a node management LIF. Using the cluster management LIF is recommended. You can use an IPv4 or IPv6 address.

command is the command you want to execute over RSH.

Example of an RSH request

The following example shows how the user named "joe", who has been set up with RSH access, can issue an RSH request to run the cluster show command:

```
admin_host$ rsh 10.72.137.28 -1 joe:password cluster show
Node Health Eligibility
------
node1 true true
node2 true true
2 entries were displayed.
admin_host$
```

Learn more about cluster show in the ONTAP command reference.

Use the ONTAP command-line interface

Learn about the ONTAP command-line interface

The ONTAP command-line interface (CLI) provides a command-based view of the management interface. You enter commands at the storage system prompt, and command results are displayed in text.

The CLI command prompt is represented as cluster_name::>.

If you set the privilege level (that is, the -privilege parameter of the set command) to advanced, the prompt includes an asterisk (*), for example:

cluster_name::*>

Learn more about set in the ONTAP command reference.

Learn about the different ONTAP shells for CLI commands

The cluster has three different shells for CLI commands, the *clustershell*, the *nodeshell*, and the *systemshell*. The shells are for different purposes, and they each have a different command set.

• The clustershell is the native shell that is started automatically when you log in to the cluster.

It provides all the commands you need to configure and manage the cluster. The clustershell CLI help (triggered by ? at the clustershell prompt) displays available clustershell commands. The man command in the clustershell displays the man page for the specified clustershell command (man <command_name>). Learn more about man in the ONTAP command reference.

• The nodeshell is a special shell for commands that take effect only at the node level.

The nodeshell is accessible through the system node run command. Learn more about system node run in the ONTAP command reference.

The nodeshell CLI help (triggered by ? or help at the nodeshell prompt) displays available nodeshell commands. The man command in the nodeshell displays the man page for the specified nodeshell command.

Many commonly used nodeshell commands and options are tunneled or aliased into the clustershell and can be executed also from the clustershell.

• The systemshell is a low-level shell that is used only for diagnostic and troubleshooting purposes.

The systemshell and the associated diag account are intended for low-level diagnostic purposes. Their access requires the diagnostic privilege level and is reserved only for technical support to perform troubleshooting tasks.

Learn more about the commands described in this procedure in the ONTAP command reference.

Access of nodeshell commands and options in the clustershell

Nodeshell commands and options are accessible through the nodeshell:

system node run -node nodename

Many commonly used nodeshell commands and options are tunneled or aliased into the clustershell and can be executed also from the clustershell.

Nodeshell options that are supported in the clustershell can be accessed by using the vserver options clustershell command. To see these options, you can do one of the following:

- Query the clustershell CLI with vserver options -vserver nodename_or_clustername -option-name ?
- Learn more about vserver options in the ONTAP command reference.

If you enter a nodeshell or legacy command or option in the clustershell, and the command or option has an equivalent clustershell command, ONTAP informs you of the clustershell command to use.

If you enter a nodeshell or legacy command or option that is not supported in the clustershell, ONTAP informs you of the "not supported" status for the command or option.

Display available nodeshell commands

You can obtain a list of available nodeshell commands by using the CLI help from the nodeshell.

Steps

1. To access the nodeshell, enter the following command at the clustershell's system prompt:

system node run -node {nodename|local}

local is the node you used to access the cluster.



The system node run command has an alias command, run.

2. Enter the following command in the nodeshell to see the list of available nodeshell commands:

[commandname] help

commandname is the name of the command whose availability you want to display. If you do not include *commandname*, the CLI displays all available nodeshell commands.

You enter exit or type Ctrl-d to return to the clustershell CLI.

```
Learn more about exit in the ONTAP command reference.
```

Example of displaying available nodeshell commands

The following example accesses the nodeshell of a node named node2 and displays information for the nodeshell command environment:

```
cluster1::> system node run -node node2
Type 'exit' or 'Ctrl-D' to return to the CLI
node2> environment help
Usage: environment status |
    [status] [shelf [<adapter>[.<shelf-number>]]] |
    [status] [shelf_log] |
    [status] [shelf_log] |
    [status] [shelf_stats] |
    [status] [shelf_power_status] |
    [status] [chassis [all | list-sensors | Temperature | PSU 1 |
    PSU 2 | Voltage | SYS FAN | NVRAM6-temperature-3 | NVRAM6-battery-3]]
```

How to navigate through ONTAP CLI command directories

Commands in the CLI are organized into a hierarchy by command directories. You can run commands in the hierarchy either by entering the full command path or by navigating through the directory structure.

When using the CLI, you can access a command directory by typing the directory's name at the prompt and then pressing Enter. The directory name is then included in the prompt text to indicate that you are interacting with the appropriate command directory. To move deeper into the command hierarchy, you type the name of a command subdirectory followed by pressing Enter. The subdirectory name is then included in the prompt text and the context shifts to that subdirectory.

You can navigate through several command directories by entering the entire command. For example, you can display information about disk drives by entering the storage disk show command at the prompt. You can also run the command by navigating through one command directory at a time, as shown in the following example:

cluster1::> storage
cluster1::storage> disk
cluster1::storage disk> show

Learn more about storage disk show in the ONTAP command reference.

You can abbreviate commands by entering only the minimum number of letters in a command that makes the command unique to the current directory. For example, to abbreviate the command in the previous example, you can enter st d sh. You can also use the Tab key to expand abbreviated commands and to display a command's parameters, including default parameter values.

You can use the top command to go to the top level of the command hierarchy, and the up command or ... command to go up one level in the command hierarchy.



Commands and command options preceded by an asterisk (*) in the CLI can be executed only at the advanced privilege level or higher.

Related information

- top
- up

Understand the rules to specify values in ONTAP CLI

Most commands include one or more required or optional parameters. Many parameters require you to specify a value for them. A few rules exist for specifying values in the CLI.

• A value can be a number, a Boolean specifier, a selection from an enumerated list of predefined values, or a text string.

Some parameters can accept a comma-separated list of two or more values. Comma-separated lists of values do not need to be in quotation marks (" "). Whenever you specify text, a space, or a query character (when not meant as a query or text starting with a less-than or greater-than symbol), you must enclose the entity in quotation marks.

- The CLI interprets a question mark (?) as the command to display help information for a particular command.
- Some text that you enter in the CLI, such as command names, parameters, and certain values, is not casesensitive.

For example, when you enter parameter values for the vserver cifs commands, capitalization is ignored. However, most parameter values, such as the names of nodes, storage virtual machines (SVMs), aggregates, volumes, and logical interfaces, are case-sensitive.

• If you want to clear the value of a parameter that takes a string or a list, you specify an empty set of quotation marks ("") or a dash ("-").

• The hash sign (#), also known as the pound sign, indicates a comment for a command-line input; if used, it should appear after the last parameter in a command line.

The CLI ignores the text between # and the end of the line.

In the following example, an SVM is created with a text comment. The SVM is then modified to delete the comment:

```
cluster1::> vserver create -vserver vs0 -subtype default -rootvolume
root_vs0
-aggregate aggr1 -rootvolume-security-style unix -language C.UTF-8 -is
-repository false -ipspace ipspaceA -comment "My SVM"
cluster1::> vserver modify -vserver vs0 -comment ""
```

In the following example, a command-line comment that uses the # sign indicates what the command does.

```
cluster1::> security login create -vserver vs0 -user-or-group-name new-
admin
-application ssh -authmethod password #This command creates a new user
account
```

Learn more about security login create in the ONTAP command reference.

View ONTAP command history and rerun any command from the history

Each CLI session keeps a history of all commands issued in it. You can view the command history of the session that you are currently in. You can also reissue commands.

To view the command history, you can use the history command.

To reissue a command, you can use the redo command with one of the following arguments:

· A string that matches part of a previous command

For example, if the only volume command you have run is volume show, you can use the redo volume command to reexecute the command.

• The numeric ID of a previous command, as listed by the history command

For example, you can use the redo 4 command to reissue the fourth command in the history list.

· A negative offset from the end of the history list

For example, you can use the redo -2 command to reissue the command that you ran two commands ago.

For example, to redo the command that is third from the end of the command history, you would enter the following command:

cluster1::> redo -3

Related information

- history
- redo
- volume

ONTAP keyboard shortcuts for editing CLI commands

The command at the current command prompt is the active command. Using keyboard shortcuts enables you to edit the active command quickly. These keyboard shortcuts are similar to those of the UNIX tcsh shell and the Emacs editor.

The following table lists the keyboard shortcuts for editing CLI commands. Ctrl- indicates that you press and hold the Ctrl key while typing the character specified after it. Esc- indicates that you press and release the Esc key and then type the character specified after it.

If you want to	Use one of these keyboard shortcuts
Move the cursor back by one character	Ctrl-BBack arrow
Move the cursor forward by one character	Ctrl-FForward arrow
Move the cursor back by one word	Esc-B
Move the cursor forward by one word	Esc-F
Move the cursor to the beginning of the line	Ctrl-A
Move the cursor to the end of the line	Ctrl-E
Remove the content of the command line from the beginning of the line to the cursor, and save it in the cut buffer. The cut buffer acts like temporary memory, similar to what is called a <i>clipboard</i> in some programs.	Ctrl-U
Remove the content of the command line from the cursor to the end of the line, and save it in the cut buffer	Ctrl-K

If you want to	Use one of these keyboard shortcuts	
Remove the content of the command line from the cursor to the end of the following word, and save it in the cut buffer	Esc-D	
Remove the word before the cursor, and save it in the cut buffer	Ctrl-W	
Yank the content of the cut buffer, and push it into the command line at the cursor	Ctrl-Y	
Delete the character before the cursor	Ctrl-HBackspace	
Delete the character where the cursor is	Ctrl-D	
Clear the line	Ctrl-C	
Clear the screen	Ctrl-L	
Replace the current content of the command line with the previous entry on the history list. With each repetition of the keyboard shortcut, the history cursor moves to the previous entry.	Ctrl-PEsc-PUp arrow	
Replace the current content of the command line with the next entry on the history list. With each repetition of the keyboard shortcut, the history cursor moves to the next entry.	 Ctrl-N Esc-N Down arrow	
Expand a partially entered command or list valid input from the current editing position	• Tab • Ctrl-I	
Display context-sensitive help	?	
Escape the special mapping for the question mark (?) character. For instance, to enter a question mark into a command's argument, press Esc and then the ? character.	Esc-?	
Start TTY output	Ctrl-Q	
Stop TTY output	Ctrl-S	

Understand the privilege levels for ONTAP CLI commands

ONTAP commands and parameters are defined at three privilege levels: *admin*, *advanced*, and *diagnostic*. The privilege levels reflect the skill levels required in performing the tasks.

• admin

Most commands and parameters are available at this level. They are used for common or routine tasks.

advanced

Commands and parameters at this level are used infrequently, require advanced knowledge, and can cause problems if used inappropriately.

You use advanced commands or parameters only with the advice of support personnel.

• diagnostic

Diagnostic commands and parameters are potentially disruptive. They are used only by support personnel to diagnose and fix problems.

Set the privilege level in the ONTAP CLI

You can set the privilege level in the CLI by using the set command. Changes to privilege level settings apply only to the session you are in. They are not persistent across sessions.

Steps

1. To set the privilege level in the CLI, use the set command with the -privilege parameter.

Example of setting the privilege level

The following example sets the privilege level to advanced and then to admin:

cluster1::> set -privilege advanced Warning: These advanced commands are potentially dangerous; use them only when directed to do so by NetApp personnel. Do you wish to continue? (y or n): y cluster1::*> set -privilege admin

Learn more about set in the ONTAP command reference.

Set display preferences for the ONTAP CLI

You can set display preferences for a CLI session by using the set command and rows command. The preferences you set apply only to the session you are in. They are not persistent across sessions.

About this task

You can set the following CLI display preferences:

- The privilege level of the command session
- · Whether confirmations are issued for potentially disruptive commands
- Whether show commands display all fields
- The character or characters to use as the field separator
- · The default unit when reporting data sizes
- The number of rows the screen displays in the current CLI session before the interface pauses output

If the preferred number of rows is not specified, it is automatically adjusted based on the actual height of the terminal. If the actual height is undefined, the default number of rows is 24.

- · The default storage virtual machine (SVM) or node
- · Whether a continuing command should stop if it encounters an error

Steps

1. To set CLI display preferences, use the set command.

To set the number of rows the screen displays in the current CLI session, you can also use the rows command.

Learn more about set and rows in the ONTAP command reference.

Example of setting display preferences in the CLI

The following example sets a comma to be the field separator, sets GB as the default data-size unit, and sets the number of rows to 50:

```
cluster1::> set -showseparator "," -units GB
cluster1::> rows 50
```

Related information

- show
- set
- rows

Use query operators in the ONTAP CLI

The management interface supports queries and UNIX-style patterns and wildcards to enable you to match multiple values in command-parameter arguments.

The following table describes the supported query operators:

Operator	Description
*	Wildcard that matches all entries.
	For example, the command volume show -volume *tmp* displays a list of all volumes whose names include the string tmp.
!	NOT operator.
	Indicates a value that is not to be matched; for example, !vs0 indicates not to match the value vs0.
	OR operator.
	Separates two values that are to be compared; for example, vs0 vs2 matches either vs0 or vs2. You can specify multiple OR statements; for example, $a \mid b^* \mid *c^*$ matches the entry a, any entry that starts with b, and any entry that includes c.
	Range operator.
	For example, 510 matches any value from 5 to 10, inclusive.
<	Less-than operator.
	For example, <20 matches any value that is less than 20.
>	Greater-than operator.
	For example, >5 matches any value that is greater than 5.
<=	Less-than-or-equal-to operator.
	For example, \leftarrow 5 matches any value that is less than or equal to 5.
>=	Greater-than-or-equal-to operator.
	For example, >=5 matches any value that is greater than or equal to 5.
{query}	Extended query.
	An extended query must be specified as the first argument after the command name, before any other parameters.
	For example, the command volume modify {-volume *tmp*} -state offline sets offline all volumes whose names include the string tmp.

If you want to parse query characters as literals, you must enclose the characters in double quotes (for example, "<10", "0..100", "*abc*", or "a|b") for the correct results to be returned.

You must enclose raw file names in double quotes to prevent the interpretation of special characters. This also applies to special characters used by the clustershell.

You can use multiple query operators in one command line. For example, the command volume show -size >1GB -percent-used <50 -vserver !vs1 displays all volumes that are greater than 1 GB in size, less than 50% used, and not in the storage virtual machine (SVM) named "vs1".

Related information

Keyboard shortcuts for editing CLI commands

Use extended queries with modify and delete commands in the ONTAP CLI

You can use extended queries to match and perform operations on objects that have specified values.

You specify extended queries by enclosing them within curly brackets ({}). An extended query must be specified as the first argument after the command name, before any other parameters. For example, to set offline all volumes whose names include the string tmp, you run the command in the following example:

cluster1::> volume modify {-volume *tmp*} -state offline

Extended queries are generally useful only with modify and delete commands. They have no meaning in create or show commands.

The combination of queries and modify operations is a useful tool. However, it can potentially cause confusion and errors if implemented incorrectly. For example, using the (advanced privilege) system node image modify command to set a node's default software image automatically sets the other software image not to be the default. The command in the following example is effectively a null operation:

cluster1::*> system node image modify {-isdefault true} -isdefault false

This command sets the current default image as the non-default image, then sets the new default image (the previous non-default image) to the non-default image, resulting in the original default settings being retained. To perform the operation correctly, you can use the command as given in the following example:

cluster1::*> system node image modify {-iscurrent false} -isdefault true

Limit the output of the ONTAP show command using the fields parameter

When you use the -instance parameter with a show command to display details, the output can be lengthy and include more information than you need. The -fields parameter of a show command enables you to display only the information you specify.

For example, running volume show -instance is likely to result in several screens of information. You can use volume show -fields fieldname[,fieldname...] to customize the output so that it includes only the specified field or fields (in addition to the default fields that are always displayed.) You can use -fields ? to display valid fields for a show command.

The following example shows the output difference between the -instance parameter and the -fields parameter:

```
cluster1::> volume show -instance
                              Vserver Name: cluster1-1
                               Volume Name: vol0
                            Aggregate Name: aggr0
                               Volume Size: 348.3GB
                         Volume Data Set ID: -
                   Volume Master Data Set ID: -
                              Volume State: online
                               Volume Type: RW
                              Volume Style: flex
                                    . . .
                      Space Guarantee Style: volume
                   Space Guarantee in Effect: true
                                    . . .
Press <space> to page down, <return> for next line, or 'q' to quit...
. . .
cluster1::>
cluster1::> volume show -fields space-guarantee, space-guarantee-enabled
vserver volume space-guarantee space-guarantee-enabled
cluster1-1 vol0 volume
                            true
cluster1-2 vol0 volume true
vs1 root_vol
               volume true
vs2 new_vol
               volume true
vs2 root_vol
               volume true
. . .
cluster1::>
```

Use the ONTAP CLI positional parameters in command input

You can take advantage of the positional parameter functionality of the ONTAP CLI to increase efficiency in command input. You can query a command to identify parameters that are positional for the command.

What a positional parameter is

• A positional parameter is a parameter that does not require you to specify the parameter name before

specifying the parameter value.

- A positional parameter can be interspersed with nonpositional parameters in the command input, as long as it observes its relative sequence with other positional parameters in the same command, as indicated in the *command_name* ? output.
- A positional parameter can be a required or optional parameter for a command.
- A parameter can be positional for one command but nonpositional for another.



Using the positional parameter functionality in scripts is not recommended, especially when the positional parameters are optional for the command or have optional parameters listed before them.

Identify a positional parameter

You can identify a positional parameter in the *command_name* ? command output. A positional parameter has square brackets surrounding its parameter name, in one of the following formats:

- [-parameter name] parameter value shows a required parameter that is positional.
- [[-parameter_name] parameter_value] shows an optional parameter that is positional.

For example, when displayed as the following in the *command_name* ? output, the parameter is positional for the command it appears in:

- [-lif] <lif-name>
- [[-lif] <lif-name>]

However, when displayed as the following, the parameter is nonpositional for the command it appears in:

- -lif <lif-name>
- [-lif <lif-name>]

Examples of using positional parameters

In the following example, the *volume create* ? output shows that three parameters are positional for the command: -volume, -aggregate, and -size.

```
cluster1::> volume create ?
   -vserver <vserver name>
                                           Vserver Name
   [-volume] <volume name>
                                           Volume Name
   [-aggregate] <aggregate name>
                                          Aggregate Name
  [[-size] {<integer>[KB|MB|GB|TB|PB]}]
                                         Volume Size
  [ -state {online|restricted|offline|force-online|force-offline|mixed} ]
                                           Volume State (default: online)
  [ -type {RW|DP|DC} ]
                                           Volume Type (default: RW)
  [ -policy <text> ]
                                           Export Policy
  [ -user <user name> ]
                                           User ID
  . . .
  [ -space-guarantee|-s {none|volume} ] Space Guarantee Style (default:
volume)
  [ -percent-snapshot-space <percent> ] Space Reserved for Snapshot
Copies
  . . .
```

In the following example, the volume create command is specified without taking advantage of the positional parameter functionality:

```
cluster1::> volume create -vserver svm1 -volume vol1 -aggregate aggr1 -size 1g
-percent-snapshot-space 0
```

The following examples use the positional parameter functionality to increase the efficiency of the command input. The positional parameters are interspersed with nonpositional parameters in the volume create command, and the positional parameter values are specified without the parameter names. The positional parameters are specified in the same sequence indicated by the volume create ? output. That is, the value for -volume is specified before that of -aggregate, which is in turn specified before that of -size.

```
cluster1::> volume create vol2 aggr1 1g -vserver svm1 -percent-snapshot-space 0
cluster1::> volume create -vserver svm1 vol3 -snapshot-policy default aggr1
-nvfail off 1g -space-guarantee none
```

How to acess the ONTAP CLI man pages

ONTAP manual (man) pages explain how to use ONTAP CLI commands. These pages are available at the command line and are also published in release-specific *command references*.

At the ONTAP command line, use the man <command_name> command to display the manual page of the specified command. If you do not specify a command name, the manual page index is displayed. You can use the man man command to view information about the man command itself. You can exit a man page by entering **q**.

Learn more about the admin-level and advanced-level ONTAP commands available in your release in the ONTAP command reference.

Record an ONTAP CLI session and manage the recorded sessions

You can record a CLI session into a file with a specified name and size limit, then upload the file to an FTP or HTTP destination. You can also display or delete files in which you previously recorded CLI sessions.

Record a CLI session

A record of a CLI session ends when you stop the recording or end the CLI session, or when the file reaches the specified size limit. The default file size limit is 1 MB. The maximum file size limit is 2 GB.

Recording a CLI session is useful, for example, if you are troubleshooting an issue and want to save detailed information or if you want to create a permanent record of space usage at a specific point in time.

Steps

1. Start recording the current CLI session into a file:

system script start

Learn more about system script start in the ONTAP command reference.

ONTAP starts recording your CLI session into the specified file.

- 2. Proceed with your CLI session.
- 3. When finished, stop recording the session:

system script stop

Learn more about system script stop in the ONTAP command reference.

ONTAP stops recording your CLI session.

Commands for managing records of CLI sessions

You use the system script commands to manage records of CLI sessions.

If you want to	Use this command
Start recording the current CLI session in to a specified file	system script start
Stop recording the current CLI session	system script stop
Display information about records of CLI sessions	system script show

If you want to	Use this command
Upload a record of a CLI session to an FTP or HTTP destination	system script upload
Delete a record of a CLI session	system script delete

Related information

ONTAP command reference

Commands for managing the automatic timeout period of CLI sessions

The timeout value specifies how long a CLI session remains idle before being automatically terminated. The CLI timeout value is cluster-wide. That is, every node in a cluster uses the same CLI timeout value.

By default, the automatic timeout period of CLI sessions is 30 minutes.

You use the system timeout commands to manage the automatic timeout period of CLI sessions.

If you want to	Use this command
Display the automatic timeout period for CLI sessions	system timeout show
Modify the automatic timeout period for CLI sessions	system timeout modify

Related information

ONTAP command reference

Cluster management (cluster administrators only)

View node-level details in an ONTAP cluster

You can view node names, whether the nodes are healthy, and whether they are eligible to participate in the cluster. At the advanced privilege level, you can also display whether a node holds epsilon.

Steps

1. To view information about the nodes in a cluster, use the cluster show command.

If you want the output to show whether a node holds epsilon, run the command at the advanced privilege level.

Learn more about cluster show in the ONTAP command reference.

Examples of displaying the nodes in a cluster

The following example displays information about all nodes in a four-node cluster:

cluster1::> cluster :	show		
Node	Health	Eligibility	
nodel	true	true	
node2	true	true	
node3	true	true	
node4	true	true	

The following example displays detailed information about the node named "node1" at the advanced privilege level:

```
cluster1::> set -privilege advanced
Warning: These advanced commands are potentially dangerous; use them only
when directed to do so by support personnel.
Do you want to continue? {y|n}: y
cluster1::*> cluster show -node node1
Node: node1
Node UUID: a67f9f34-9d8f-11da-b484-000423b6f094
Epsilon: false
Eligibility: true
Health: true
```

View ONTAP cluster-level details

You can display a cluster's unique identifier (UUID), name, serial number, location, and contact information.

Steps

1. To display a cluster's attributes, use the cluster identity show command.

Example of displaying cluster attributes

The following example displays the name, serial number, location, and contact information of a cluster.

```
cluster1::> cluster identity show
        Cluster UUID: 1cd8a442-86d1-11e0-ae1c-123478563412
        Cluster Name: cluster1
        Cluster Serial Number: 1-80-123456
        Cluster Location: Sunnyvale
        Cluster Contact: jsmith@example.com
```

Learn more about cluster identity show in the ONTAP command reference.

Modify ONTAP cluster attributes

You can modify a cluster's attributes, such as the cluster name, location, and contact information as needed.

About this task

You cannot change a cluster's UUID, which is set when the cluster is created.

Steps

1. To modify cluster attributes, use the cluster identity modify command.

The -name parameter specifies the name of the cluster. Learn more about cluster identity modify and the rules for specifying the cluster's name in the ONTAP command reference.

The -location parameter specifies the location for the cluster.

The -contact parameter specifies the contact information such as a name or e-mail address.

Example of renaming a cluster

The following command renames the current cluster ("cluster1") to "cluster2":

cluster1::> cluster identity modify -name cluster2

View the replication status of the ONTAP cluster rings

You can view the status of cluster replication rings to help you diagnose cluster-wide problems. If your cluster is experiencing problems, support personnel might ask you to perform this task to assist with troubleshooting efforts.

Steps

1. To display the status of cluster replication rings, use the cluster ring show command at the advanced privilege level.

Example of displaying cluster ring-replication status

The following example displays the status of the VLDB replication ring on a node named node0:

```
cluster1::> set -privilege advanced
Warning: These advanced commands are potentially dangerous; use them only
when directed to do so by support personnel.
Do you wish to continue? (y or n): y
cluster1::*> cluster ring show -node node0 -unitname vldb
        Node: node0
        Unit Name: vldb
        Status: master
        Epoch: 5
        Master Node: node0
        Local Node: node0
        DB Epoch: 5
DB Transaction: 56
        Number Online: 4
        RDB UUID: e492d2c1-fc50-11e1-bae3-123478563412
```

Learn more about cluster ring show in the ONTAP command reference.

ONTAP cluster health assessments using quorum and epsilon

Quorum and epsilon are important measures of cluster health and function that together indicate how clusters address potential communications and connectivity challenges.

Quorum is a precondition for a fully functioning cluster. When a cluster is in quorum, a simple majority of nodes are healthy and can communicate with each other. When quorum is lost, the cluster loses the ability to accomplish normal cluster operations. Only one collection of nodes can have quorum at any one time because all of the nodes collectively share a single view of the data. Therefore, if two non-communicating nodes are permitted to modify the data in divergent ways, it is no longer possible to reconcile the data into a single data view.

Each node in the cluster participates in a voting protocol that elects one node *master*; each remaining node is a *secondary*. The master node is responsible for synchronizing information across the cluster. When quorum is formed, it is maintained by continual voting. If the master node goes offline and the cluster is still in quorum, a new master is elected by the nodes that remain online.

Because there is the possibility of a tie in a cluster that has an even number of nodes, one node has an extra fractional voting weight called *epsilon*. If the connectivity between two equal portions of a large cluster fails, the group of nodes containing epsilon maintains quorum, assuming that all of the nodes are healthy. For example, the following illustration shows a four-node cluster in which two of the nodes have failed. However, because one of the surviving nodes holds epsilon, the cluster remains in quorum even though there is not a simple majority of healthy nodes.



Epsilon is automatically assigned to the first node when the cluster is created. If the node that holds epsilon becomes unhealthy, takes over its high-availability partner, or is taken over by its high-availability partner, then epsilon is automatically reassigned to a healthy node in a different HA pair.

Taking a node offline can affect the ability of the cluster to remain in quorum. Therefore, ONTAP issues a warning message if you attempt an operation that will either take the cluster out of quorum or else put it one outage away from a loss of quorum. You can disable the quorum warning messages by using the cluster quorum-service options modify command at the advanced privilege level. Learn more about cluster quorum-service options modify in the ONTAP command reference.

In general, assuming reliable connectivity among the nodes of the cluster, a larger cluster is more stable than a smaller cluster. The quorum requirement of a simple majority of half the nodes plus epsilon is easier to maintain in a cluster of 24 nodes than in a cluster of two nodes.

A two-node cluster presents some unique challenges for maintaining quorum. Two-node clusters use *cluster HA*, in which neither node holds epsilon; instead, both nodes are continuously polled to ensure that if one node fails, the other has full read-write access to data, as well as access to logical interfaces and management functions.

View storage capacity utilization of system volumes in an ONTAP cluster

System volumes are FlexVol volumes that contain special metadata, such as metadata for file services audit logs. These volumes are visible in the cluster so that you can fully account for storage use in your cluster.

System volumes are owned by the cluster management server (also called the admin SVM), and they are created automatically when file services auditing is enabled.

You can view system volumes by using the volume show command, but most other volume operations are not permitted. For example, you cannot modify a system volume by using the volume modify command.

This example shows four system volumes on the admin SVM, which were automatically created when file services auditing was enabled for a data SVM in the cluster:

cluster1::> volume show -vserver cluster1 Vserver Volume Aggregate State Туре Size Available Used% _____ _ _____ ____ cluster1 MDV aud 1d0131843d4811e296fc123478563412 aggr0 online RW 2GB 1.90GB 5% cluster1 MDV aud 8be27f813d7311e296fc123478563412 root vs0 online RW 2GB 1.90GB 5% cluster1 MDV aud 9dc4ad503d7311e296fc123478563412 online RW 2GB 1.90GB aggr1 5% cluster1 MDV aud a4b887ac3d7311e296fc123478563412 aqqr2 online RW 2GB 1.90GB 5% 4 entries were displayed.

Manage nodes

Add nodes to an ONTAP cluster

After a cluster is created, you can expand it by adding nodes to it. You add only one node at a time.

Before you begin

- If you are adding nodes to a multiple-node cluster, all the existing nodes in the cluster must be healthy (indicated by cluster show). Learn more about cluster show in the ONTAP command reference.
- If you are adding nodes to a two-node switchless cluster, you must convert your two-node switchless cluster to a switch-attached cluster using a NetApp supported cluster switch.

The switchless cluster functionality is supported only in a two-node cluster.

- If you are adding a second node to a single-node cluster, the second node must have been installed, and the cluster network must have been configured.
- If the cluster has SP automatic configuration enabled, the subnet specified for the SP must have available resources to allow the joining node to use the specified subnet to automatically configure the SP.
- You must have gathered the following information for the new node's node management LIF:
 - Port
 - IP address
 - Netmask
 - Default gateway

About this task

Nodes must be in even numbers so that they can form HA pairs. After you start to add a node to the cluster, you must complete the process. The node must be part of the cluster before you can start to add another node.

Steps

1. Power on the node that you want to add to the cluster.

The node boots, and the Node Setup wizard starts on the console.

```
Welcome to node setup.
You can enter the following commands at any time:
    "help" or "?" - if you want to have a question clarified,
    "back" - if you want to change previously answered questions, and
    "exit" or "quit" - if you want to quit the setup wizard.
    Any changes you made before quitting will be saved.
To accept a default or omit a question, do not enter a value.
Enter the node management interface port [eOM]:
```

2. Exit the Node Setup wizard: exit

The Node Setup wizard exits, and a login prompt appears, warning that you have not completed the setup tasks.

Learn more about exit in the ONTAP command reference.

- 3. Log in to the admin account by using the admin user name.
- 4. Start the Cluster Setup wizard:

```
::> cluster setup
Welcome to the cluster setup wizard.
You can enter the following commands at any time:
  "help" or "?" - if you want to have a question clarified,
  "back" - if you want to change previously answered questions, and
  "exit" or "quit" - if you want to quit the cluster setup wizard.
Any changes you made before quitting will be saved.
You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value....
Use your web browser to complete cluster setup by accessing
https://<node_mgmt_or_eOM_IP_address>
Otherwise, press Enter to complete cluster setup using the
command line interface:
```



For more information on setting up a cluster using the setup GUI, see the node management documentation. Learn more about cluster setup in the ONTAP command reference.

5. Press Enter to use the CLI to complete this task. When prompted to create a new cluster or join an existing one, enter join.

```
Do you want to create a new cluster or join an existing cluster? {create, join}: join
```

If the ONTAP version running on the new node is different to the version running on the existing cluster, the system reports a System checks Error: Cluster join operation cannot be performed at this time error. This is the expected behavior. To continue, run the cluster add-node -allow -mixed-version-join true -cluster-ips <IP address> -node-names <new_node_name> command at the advanced privilege level from an existing node in the cluster.

- 6. Follow the prompts to set up the node and join it to the cluster:
 - To accept the default value for a prompt, press Enter.
 - To enter your own value for a prompt, enter the value, and then press Enter.
- 7. Repeat the preceding steps for each additional node that you want to add.

After you finish

After adding nodes to the cluster, you should enable storage failover for each HA pair.

Related information

• Mixed version clusters supported for ONTAP software upgrades

• cluster add-node

Remove nodes from an ONTAP cluster

You can remove unwanted nodes from a cluster, one node at a time. After you remove a node, you must also remove its failover partner. If you are removing a node, then its data becomes inaccessible or erased.

Before you begin

The following conditions must be satisfied before removing nodes from the cluster:

- More than half of the nodes in the cluster must be healthy.
- All data, volumes, and non-root aggregates have been relocated or removed from the node.
 - All of the data on the node that you want to remove must have been evacuated. This might include purging data from an encrypted volume.
 - All non-root volumes have been moved from aggregates owned by the node.
 - All non-root aggregates have been deleted from the node.
- All LIFs and VLANs have been relocated or removed from the node.
 - Data LIFs have been deleted or relocated from the node.
 - Cluster management LIFs have been relocated from the node and the home ports changed.
 - All intercluster LIFs have been removed. When you remove intercluster LIFs a warning is displayed that can be ignored.
 - All VLANs on the node have been deleted.
- The node is not participating in any failover relationships.
 - Storage failover has been disabled for the node.
 - All LIF failover rules have been modified to remove ports on the node.
- If the node owns Federal Information Processing Standards (FIPS) disks or self-encrypting disks (SEDs), disk encryption has been removed by returning the disks to unprotected mode.
 - You might also want to sanitize FIPS drives or SEDs.
- If you have LUNs on the node to be removed, you should modify the Selective LUN Map (SLM) reportingnodes list before you remove the node.

If you do not remove the node and its HA partner from the SLM reporting-nodes list, access to the LUNs previously on the node can be lost even though the volumes containing the LUNs were moved to another node.

It is recommended that you issue an AutoSupport message to notify NetApp technical support that node removal is underway.



Do not perform operations such as cluster remove-node, cluster unjoin, and node rename when an automated ONTAP upgrade is in progress.

About this task

• If you are running a mixed-version cluster, you can remove the last low-version node by using one of the advanced privilege commands beginning with ONTAP 9.3:

- ONTAP 9.3: cluster unjoin -skip-last-low-version-node-check
- ONTAP 9.4 and later: cluster remove-node -skip-last-low-version-node-check
- If you unjoin 2 nodes from a 4-node cluster, cluster HA is automatically enabled on the two remaining nodes.



All system and user data, from all disks that are connected to the node, must be made inaccessible to users before removing a node from the cluster. If a node was incorrectly unjoined from a cluster, contact NetApp Support for assistance with options for recovery.

Steps

1. Change the privilege level to advanced:

set -privilege advanced

2. Verify if a node on the cluster holds epsilon:

```
cluster show -epsilon true
```

- If a node on the cluster holds epsilon and that node is going to be unjoined, move epsilon to a node that is not going to be unjoined:
 - a. Move epsilon from the node that is going to be unjoined

cluster modify -node <name of node to be unjoined> -epsilon false

b. Move epsilon to a node that is not going to be unjoined:

cluster modify -node <node name> -epsilon true

4. Identify the current master node:

cluster ring show

The master node is the node that holds processes such as mgmt, vldb, vifmgr, bcomd, and crs.

- 5. If the node you want to remove is the current master node, then enable another node in the cluster to be elected as the master node:
 - a. Make the current master node ineligible to participate in the cluster:

cluster modify -node <node name> -eligibility false

This will cause the node to be marked as unhealthy until eligibility is restored in the next step. When

the master node become ineligible, one of the remaining nodes is elected by the cluster quorum as the new master.

b. Make the previous master node eligible to participate in the cluster again:

```
cluster modify -node <node name> -eligibility true
```

- 6. Log into the remote node management LIF or the cluster-management LIF on a node other than the one that is being removed.
- 7. Remove the node from the cluster:

For this ONTAP version	Use this command	
ONTAP 9.3	cluster unjoin	
ONTAP 9.4 and later	With node name:	
	cluster remove-node -node <node_name></node_name>	
	With node IP:	
	cluster remove-node -cluster_ip <node_ip></node_ip>	

If you have a mixed version cluster and you are removing the last lower version node, use the <code>-skip -last-low-version-node-check</code> parameter with these commands.

The system informs you of the following:

- You must also remove the node's failover partner from the cluster.
- After the node is removed and before it can rejoin a cluster, you must use boot menu option (4) Clean configuration and initialize all disks or option (9) Configure Advanced Drive Partitioning to erase the node's configuration and initialize all disks.

A failure message is generated if you have conditions that you must address before removing the node. For example, the message might indicate that the node has shared resources that you must remove or that the node is in a cluster HA configuration or storage failover configuration that you must disable.

If the node is the quorum master, the cluster will briefly lose and then return to quorum. This quorum loss is temporary and does not affect any data operations.

8. If a failure message indicates error conditions, address those conditions and rerun the cluster removenode or cluster unjoin command. The node is automatically rebooted after it is successfully removed from the cluster.

- 9. If you are repurposing the node, erase the node configuration and initialize all disks:
 - a. During the boot process, press Ctrl-C to display the boot menu when prompted to do so.
 - b. Select the boot menu option (4) Clean configuration and initialize all disks.
- 10. Return to admin privilege level:

set -privilege admin

11. Repeat the preceding steps to remove the failover partner from the cluster.

Related information

• cluster remove-node

Access ONTAP node logs, core dumps, and MIB files using a web browser

The Service Processor Infrastructure (spi) web service is enabled by default to enable a web browser to access the log, core dump, and MIB files of a node in the cluster. The files remain accessible even when the node is down, provided that the node is taken over by its partner.

Before you begin

• The cluster management LIF must be up.

You can use the management LIF of the cluster or a node to access the spi web service. However, using the cluster management LIF is recommended.

The network interface show command displays the status of all LIFs in the cluster.

Learn more about network interface show in the ONTAP command reference.

- You must use a local user account to access the spi web service, domain user accounts are not supported.
- If your user account does not have the admin role (which has access to the spi web service by default), your access-control role must be granted access to the spi web service.

The vserver services web access show command shows what roles are granted access to which web services.

• If you are not using the admin user account (which includes the http access method by default), your user account must be set up with the http access method.

The security login show command shows user accounts' access and login methods and their access-control roles.

Learn more about security login show in the ONTAP command reference.

• If you want to use HTTPS for secure web access, SSL must be enabled and a digital certificate must be installed.

The system services web show command displays the configuration of the web protocol engine at the cluster level.

About this task

The spi web service is enabled by default, and the service can be disabled manually (vserver services web modify -vserver * -name spi -enabled false).

The admin role is granted access to the spi web service by default, and the access can be disabled manually (services web access delete -vserver *cluster name* -name spi -role admin).

Steps

1. Point the web browser to the spi web service URL in one of the following formats:

```
° http://cluster-mgmt-LIF/spi/
```

° https://cluster-mgmt-LIF/spi/

cluster-mgmt-LIF is the IP address of the cluster management LIF.

2. When prompted by the browser, enter your user account and password.

After your account is authenticated, the browser displays links to the /mroot/etc/log/, /mroot/etc/crash/, and /mroot/etc/mib/ directories of each node in the cluster.

Access the system console of an ONTAP node

If a node is hanging at the boot menu or the boot environment prompt, you can access it only through the system console (also called the *serial console*). You can access the system console of a node from an SSH connection to the node's SP or to the cluster.

About this task

Both the SP and ONTAP offer commands that enable you to access the system console. However, from the SP, you can access only the system console of its own node. From the cluster, you can access the system console of any other node in the cluster (other than the local node).

Steps

1. Access the system console of a node:

If you are in the	Enter this command
SP CLI of the node	system console
ONTAP CLI	system node run-console

- 2. Log in to the system console when you are prompted to do so.
- 3. To exit the system console, press Ctrl-D.

Examples of accessing the system console

The following example shows the result of entering the system console command at the "SP node2" prompt. The system console indicates that node2 is hanging at the boot environment prompt. The

boot_ontap command is entered at the console to boot the node to ONTAP. Ctrl-D is then pressed to exit the console and return to the SP.

(Ctrl-D is pressed to exit the system console.)

```
Connection to 123.12.123.12 closed. SP node2>
```

The following example shows the result of entering the system node run-console command from ONTAP to access the system console of node2, which is hanging at the boot environment prompt. The boot_ontap command is entered at the console to boot node2 to ONTAP. Ctrl-D is then pressed to exit the console and return to ONTAP.

(Ctrl-D is pressed to exit the system console.)

Connection to 123.12.123.12 closed. cluster1::>

Manage ONTAP node root volumes and root aggregates

A node's root volume is a FlexVol volume that is installed at the factory or by setup software. It is reserved for system files, log files, and core files. The directory name is /mroot, which is accessible only through the systemshell by technical support. The minimum size for a node's root volume depends on the platform model.

Rules governing node root volumes and root aggregates overview

A node's root volume contains special directories and files for that node. The root aggregate contains the root volume. A few rules govern a node's root volume and root aggregate.

- The following rules govern the node's root volume:
 - Unless technical support instructs you to do so, do not modify the configuration or content of the root volume.
 - Do not store user data in the root volume.

Storing user data in the root volume increases the storage giveback time between nodes in an HA pair.

- You can move the root volume to another aggregate. See Relocate root volumes to new aggregates.
- The root aggregate is dedicated to the node's root volume only.

ONTAP prevents you from creating other volumes in the root aggregate.

NetApp Hardware Universe

Free up space on a node's root volume

A warning message appears when a node's root volume has become full or almost full. The node cannot operate properly when its root volume is full. You can free up space on a node's root volume by deleting core dump files, packet trace files, and root volume snapshots.

Steps

1. Display the node's core dump files and their names:

system node coredump show

2. Delete unwanted core dump files from the node:

system node coredump delete

3. Access the nodeshell:

system node run -node nodename

nodename is the name of the node whose root volume space you want to free up.
4. Switch to the nodeshell advanced privilege level from the nodeshell:

priv set advanced

- 5. Display and delete the node's packet trace files through the nodeshell:
 - a. Display all files in the node's root volume:

ls /etc

b. If any packet trace files (*.trc) are in the node's root volume, delete them individually:

rm /etc/log/packet_traces/file_name.trc

- 6. Identify and delete the node's root volume snapshots through the nodeshell:
 - a. Identify the root volume name:

vol status

The root volume is indicated by the word "root" in the "Options" column of the vol status command output.

In the following example, the root volume is vol0:

```
node1*> vol status
Volume State
vol0 online
```

Status raid_dp, flex 64-bit

Options root, nvfail=on

b. Display root volume snapshots:

snap list root_vol_name

- c. Delete unwanted root volume snapshots:
 - snap delete root_vol_namesnapshot_name
- 7. Exit the nodeshell and return to the clustershell:

exit

Relocate root volumes to new aggregates

The root replacement procedure migrates the current root aggregate to another set of disks without disruption.

About this task

Storage failover must be enabled to relocate root volumes. You can use the storage failover modify -node *nodename* -enable true command to enable failover.

You can change the location of the root volume to a new aggregate in the following scenarios:

- When the root aggregates are not on the disk you prefer
- When you want to rearrange the disks connected to the node
- When you are performing a shelf replacement of the EOS disk shelves

Steps

1. Set the privilege level to advanced:

set privilege advanced

2. Relocate the root aggregate:

system node migrate-root -node *nodename* -disklist *disklist* -raid-type *raid-type*

• -node

Specifies the node that owns the root aggregate that you want to migrate.

• -disklist

Specifies the list of disks on which the new root aggregate will be created. All disks must be spares and owned by the same node. The minimum number of disks required is dependent on the RAID type.

• -raid-type

Specifies the RAID type of the root aggregate. The default value is raid-dp.

3. Monitor the progress of the job:

job show -id jobid -instance

Results

If all of the pre-checks are successful, the command starts a root volume replacement job and exits. Expect the node to restart.

Start or stop an ONTAP node for maintenance or troubleshooting

You might need to start or stop a node for maintenance or troubleshooting reasons. You can do so from the ONTAP CLI, the boot environment prompt, or the SP CLI.

Using the SP CLI command system power off or system power cycle to turn off or power-cycle a node might cause an improper shutdown of the node (also called a *dirty shutdown*) and is not a substitute for a graceful shutdown using the ONTAP system node halt command.

Reboot a node at the system prompt

You can reboot a node in normal mode from the system prompt. A node is configured to boot from the boot device, such as a PC CompactFlash card.

Steps

1. If the cluster contains four or more nodes, verify that the node to be rebooted does not hold epsilon:

a. Set the privilege level to advanced:

set -privilege advanced

b. Determine which node holds epsilon:

cluster show

The following example shows that "node1" holds epsilon:

```
cluster1::*> cluster show
Node
             Health Eligibility Epsilon
----- -----
                            _____
node1
             true true
                             true
node2
                            false
             true true
node3
             true true
                            false
           true true
node4
                            false
4 entries were displayed.
```

c. If the node to be rebooted holds epsilon, then remove epsilon from the node:

cluster modify -node node name -epsilon false

d. Assign epsilon to a different node that will remain up:

```
cluster modify -node node_name -epsilon true
```

e. Return to the admin privilege level:

set -privilege admin

2. Use the system node reboot command to reboot the node.

If you do not specify the <code>-skip-lif-migration</code> parameter, the command attempts to migrate data and cluster management LIFs synchronously to another node prior to the reboot. If the LIF migration fails or times out, the rebooting process is aborted, and ONTAP displays an error to indicate the LIF migration failure.

cluster1::> system node reboot -node node1 -reason "software upgrade"

The node begins the reboot process. The ONTAP login prompt appears, indicating that the reboot process is complete.

Boot ONTAP at the boot environment prompt

You can boot the current release or the backup release of ONTAP when you are at the boot environment prompt of a node.

Steps

1. Access the boot environment prompt from the storage system prompt by using the system node halt command.

The storage system console displays the boot environment prompt.

2. At the boot environment prompt, enter one of the following commands:

To boot	Enter
The current release of ONTAP	boot_ontap
The ONTAP primary image from the boot device	boot_primary
The ONTAP backup image from the boot device	boot_backup

If you are unsure about which image to use, you should use boot ontap in the first instance.

Shut down a node

You can shut down a node if it becomes unresponsive or if support personnel direct you to do so as part of troubleshooting efforts.

Steps

- 1. If the cluster contains four or more nodes, verify that the node to be shut down does not hold epsilon:
 - a. Set the privilege level to advanced:

set -privilege advanced

b. Determine which node holds epsilon:

cluster show

The following example shows that "node1" holds epsilon:

```
cluster1::*> cluster show
Node
               Health Eligibility Epsilon
_____
node1
               true
                    true
                               true
node2
               true true
                               false
node3
                              false
               true
                    true
node4
               true
                    true
                               false
4 entries were displayed.
```

c. If the node to be shut down holds epsilon, then remove epsilon from the node:

cluster modify -node node_name -epsilon false

d. Assign epsilon to a different node that will remain up:

cluster modify -node node_name -epsilon true

e. Return to the admin privilege level:

set -privilege admin

2. Use the system node halt command to shut down the node.

If you do not specify the <code>-skip-lif-migration</code> parameter, the command attempts to migrate data and cluster management LIFs synchronously to another node prior to the shutdown. If the LIF migration fails or times out, the shutdown process is aborted, and ONTAP displays an error to indicate the LIF migration failure.

You can manually trigger a core dump with the shutdown by using both the -dump parameter.

The following example shuts down the node named "node1" for hardware maintenance:

cluster1::> system node halt -node node1 -reason 'hardware maintenance'

Manage an ONTAP node using the boot menu

You can use the boot menu to correct configuration problems on a node, reset the admin password, initialize disks, reset the node configuration, and restore the node configuration information back to the boot device.



If an HA pair is using encrypting SAS or NVMe drives (SED, NSE, FIPS), you must follow the instructions in the topic Returning a FIPS drive or SED to unprotected mode for all drives within the HA pair prior to initializing the system (boot options 4 or 9). Failure to do this may result in future data loss if the drives are repurposed.

Steps

1. Reboot the node to access the boot menu by using the system node reboot command at the system prompt.

The node begins the reboot process.

2. During the reboot process, press Ctrl-C to display the boot menu when prompted to do so.

The node displays the following options for the boot menu:

```
    Normal Boot.
    Boot without /etc/rc.
    Change password.
    Clean configuration and initialize all disks.
    Maintenance mode boot.
    Update flash from backup config.
    Install new software first.
    Reboot node.
    Configure Advanced Drive Partitioning.
    Set onboard key management recovery secrets.
    Configure node for external key management.
```



Boot menu option (2) Boot without /etc/rc is obsolete and takes no effect on the system.

3. Select one of the following options by entering the corresponding number:

То	Select	
Continue to boot the node in normal mode	1) Normal Boot	
Change the password of the node, which is also the "admin" account password	3) Change Password	
Initialize the node's disks and create a root volume for the node	 4) Clean configuration and initialize all disks i This menu option erases all data on the disks of the node and resets your node configuration to the factory default settings. Only select this menu item after the node has been removed from a cluster (unjoined) and is not joined to another cluster. For a node with internal or external disk shelves, the root volume on the internal disks is initialized. If there are no internal disk shelves, then the root volume on the external disks is initialized. If the node you want to initialize has disks that are partitioned for root-data partitioning, the disks must be unpartitioned before the node can be initialized, see 9) Configure Advanced Drive Partitioning and Disks and aggregates management. 	
Perform aggregate and disk maintenance operations and obtain detailed aggregate and disk information.	5) Maintenance mode boot You exit Maintenance mode by using the halt command.	

То	Select
Restore the configuration information from the node's root volume to the boot device, such as a PC CompactFlash card	6) Update flash from backup config ONTAP stores some node configuration information on the boot device. When the node reboots, the information on the boot device is automatically backed up onto the node's root volume. If the boot device becomes corrupted or needs to be replaced, you must use this menu option to restore the configuration information from the node's root volume back to the boot device.
Install new software on the node	 7) Install new software first If the ONTAP software on the boot device does not include support for the storage array that you want to use for the root volume, you can use this menu option to obtain a version of the software that supports your storage array and install it on the node. This menu option is only for installing a newer version of ONTAP software on a node that has no root volume installed. Do <i>not</i> use this menu option to upgrade ONTAP.
Reboot the node	8) Reboot node
Unpartition all disks and remove their ownership information or clean the configuration and initialize the system with whole or partitioned disks	 9) Configure Advanced Drive Partitioning The Advanced Drive Partitioning option provides additional management features for disks that are configured for root-data or root-data-data partitioning. The following options are available from Boot Option 9: (9a) Unpartition all disks and remove their
	ownership information. (9b) Clean configuration and initialize system with partitioned disks. (9c) Clean configuration and initialize system with whole disks. (9d) Reboot the node. (9e) Return to main boot menu.

View the attributes of the nodes in an ONTAP cluster

You can view the attributes of one or more nodes in the cluster, for example, the name, owner, location, model number, serial number, how long the node has been running, health state, and eligibility to participate in a cluster.

Steps

1. To display the attributes of a specified node or about all nodes in a cluster, use the system node show

Example of displaying information about a node

The following example displays detailed information about node1:

```
cluster1::> system node show -node node1
                          Node: node1
                         Owner: Eng IT
                      Location: Lab 5
                         Model: model number
                 Serial Number: 12345678
                     Asset Tag: -
                        Uptime: 23 days 04:42
               NVRAM System ID: 118051205
                     System ID: 0118051205
                        Vendor: NetApp
                        Health: true
                   Eligibility: true
       Differentiated Services: false
           All-Flash Optimized: true
            Capacity Optimized: false
                 QLC Optimized: false
   All-Flash Select Optimized: false
SAS2/SAS3 Mixed Stack Support: none
```

Modify the attributes of an ONTAP node

You can modify the attributes of a node as required. The attributes that you can modify include the node's owner information, location information, asset tag, and eligibility to participate in the cluster.

About this task

A node's eligibility to participate in the cluster can be modified at the advanced privilege level by using the -eligibility parameter of the system node modify or cluster modify command. If you set a node's eligibility to false, the node becomes inactive in the cluster.



You cannot modify node eligibility locally. It must be modified from a different node. Node eligiblity also cannot be modified with a cluster HA configuration.



You should avoid setting a node's eligibility to false, except for situations such as restoring the node configuration or prolonged node maintenance. SAN and NAS data access to the node might be impacted when the node is ineligible.

Steps

1. Use the system node modify command to modify a node's attributes.

Example of modifying node attributes

The following command modifies the attributes of the "node1" node. The node's owner is set to "Joe Smith" and its asset tag is set to "js1234":

```
cluster1::> system node modify -node node1 -owner "Joe Smith" -assettag
js1234
```

Related information

- system node modify
- · cluster modify

Rename an ONTAP node

You can change a node's name as required.

Steps

1. To rename a node, use the system node rename command.

The -newname parameter specifies the new name for the node. Learn more about system node rename in the ONTAP command reference.

If you want to rename multiple nodes in the cluster, you must run the command for each node individually.



Node name cannot be "all" because "all" is a system reserved name.

Example of renaming a node

The following command renames node "node1" to "node1a":

cluster1::> system node rename -node node1 -newname node1a

Manage a single-node ONTAP cluster

A single-node cluster is a special implementation of a cluster running on a standalone node. Single-node clusters are not recommended because they do not provide redundancy. If the node goes down, data access is lost.



For fault tolerance and nondisruptive operations, it is highly recommended that you configure your cluster with high-availability (HA pairs).

If you choose to configure or upgrade a single-node cluster, you should be aware of the following:

- Root volume encryption is not supported on single-node clusters.
- If you remove nodes to have a single-node cluster, you should modify the cluster ports to serve data traffic by modifying the cluster ports to be data ports, and then creating data LIFs on the data ports.
- For single-node clusters, you can specify the configuration backup destination during software setup. After setup, those settings can be modified using ONTAP commands.

• If there are multiple hosts connecting to the node, each host can be configured with a different operating system such as Windows or Linux. If there are multiple paths from the host to the controller, then ALUA must be enabled on the host.

Ways to configure iSCSI SAN hosts with single nodes

You can configure iSCSI SAN hosts to connect directly to a single node or to connect through one or more IP switches. The node can have multiple iSCSI connections to the switch.

Direct-attached single-node configurations

In direct-attached single-node configurations, one or more hosts are directly connected to the node.



Single-network single-node configurations

In single-network single-node configurations, one switch connects a single node to one or more hosts. Because there is a single switch, this configuration is not fully redundant.



Multi-network single-node configurations

In multi-network single-node configurations, two or more switches connect a single node to one or more hosts. Because there are multiple switches, this configuration is fully redundant.



Ways to configure FC and FC-NVMe SAN hosts with single nodes

You can configure FC and FC-NVMe SAN hosts with single nodes through one or more fabrics. N-Port ID Virtualization (NPIV) is required and must be enabled on all FC switches in the fabric. You cannot directly attach FC or FC-NMVE SAN hosts to single nodes without using an FC switch.

Single-fabric single-node configurations

In single-fabric single-node configurations, there is one switch connecting a single node to one or more hosts. Because there is a single switch, this configuration is not fully redundant.

In single-fabric single-node configurations, multipathing software is not required if you only have a single path from the host to the node.

Multifabric single-node configurations

In multifabric single-node configurations, there are two or more switches connecting a single node to one or more hosts. For simplicity, the following figure shows a multifabric single-node configuration with only two fabrics, but you can have two or more fabrics in any multifabric configuration. In this figure, the storage controller is mounted in the top chassis and the bottom chassis can be empty or can have an IOMX module, as it does in this example.

The FC target ports (0a, 0c, 0b, 0d) in the illustrations are examples. The actual port numbers vary depending on the model of your storage node and whether you are using expansion adapters.



Related information

NetApp Technical Report 4684: Implementing and Configuring Modern SANs with NVMe-oF

ONTAP upgrade for single-node cluster

You can use the ONTAP CLI to perform an automated update of a single-node cluster. Single-node clusters lack redundancy, this means that updates are always disruptive. You can't perform disruptive upgrades using System Manager.

Before you begin

You must complete upgrade preparation steps.

Steps

1. Delete the previous ONTAP software package:

```
cluster image package delete -version <previous package version>
```

2. Download the target ONTAP software package:

```
cluster image package get -url location
```

```
cluster1::> cluster image package get -url
http://www.example.com/software/9.7/image.tgz
```

```
Package download completed.
Package processing completed.
```

3. Verify that the software package is available in the cluster package repository:

cluster image package show-repository

cluster1::> cluster image package show-repository
Package Version Package Build Time
-----9.7 M/DD/YYYY 10:32:15

4. Verify that the cluster is ready to be upgraded:

cluster image validate -version <package_version_number>

cluster1::> cluster image validate -version 9.7

WARNING: There are additional manual upgrade validation checks that must be performed after these automated validation checks have completed...

5. Monitor the progress of the validation:

cluster image show-update-progress

- 6. Complete all required actions identified by the validation.
- 7. Optionally, generate a software upgrade estimate:

cluster image update -version <package version number> -estimate-only

The software upgrade estimate displays details about each component to be updated, and the estimated duration of the upgrade.

8. Perform the software upgrade:

cluster image update -version <package_version_number>



If an issue is encountered, the update pauses and prompts you to take corrective action. You can use the cluster image show-update-progress command to view details about any issues and the progress of the update. After correcting the issue, you can resume the update by using the cluster image resume-update command.

9. Display the cluster update progress:

cluster image show-update-progress

The node is rebooted as part of the update and cannot be accessed while rebooting.

10. Trigger a notification:

autosupport invoke -node * -type all -message "Finishing Upgrade"

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

Configure the SP/BMC network

Isolate ONTAP management traffic in the network

It is a best practice to configure SP/BMC and the e0M management interface on a subnet dedicated to management traffic. Running data traffic over the management network can cause performance degradation and routing problems.

The management Ethernet port on most storage controllers (indicated by a wrench icon on the rear of the chassis) is connected to an internal Ethernet switch. The internal switch provides connectivity to SP/BMC and to the e0M management interface, which you can use to access the storage system via TCP/IP protocols like Telnet, SSH, and SNMP.



If you plan to use both the remote management device and e0M, you must configure them on the same IP subnet. Since these are low-bandwidth interfaces, the best practice is to configure SP/BMC and e0M on a subnet dedicated to management traffic.

If you cannot isolate management traffic, or if your dedicated management network is unusually large, you should try to keep the volume of network traffic as low as possible. Excessive ingress broadcast or multicast traffic may degrade SP/BMC performance.



Some storage controllers, such as the AFF A800, have two external ports, one for BMC and the other for e0M. For these controllers, there is no requirement to configure BMC and e0M on the same IP subnet.

Learn about ONTAP SP/BMC network configuration

You can enable cluster-level, automatic network configuration for the SP (recommended). You can also leave the SP automatic network configuration disabled (the default) and manage the SP network configuration manually at the node level. A few considerations exist for each case.



This topic applies to both the SP and the BMC.

The SP automatic network configuration enables the SP to use address resources (including the IP address, subnet mask, and gateway address) from the specified subnet to set up its network automatically. With the SP automatic network configuration, you do not need to manually assign IP addresses for the SP of each node. By default, the SP automatic network configuration is disabled; this is because enabling the configuration requires that the subnet to be used for the configuration be defined in the cluster first.

If you enable the SP automatic network configuration, the following scenarios and considerations apply:

- If the SP has never been configured, the SP network is configured automatically based on the subnet specified for the SP automatic network configuration.
- If the SP was previously configured manually, or if the existing SP network configuration is based on a different subnet, the SP network of all nodes in the cluster are reconfigured based on the subnet that you specify in the SP automatic network configuration.

The reconfiguration could result in the SP being assigned a different address, which might have an impact on your DNS configuration and its ability to resolve SP host names. As a result, you might need to update your DNS configuration.

- A node that joins the cluster uses the specified subnet to configure its SP network automatically.
- The system service-processor network modify command does not enable you to change the SP IP address.

When the SP automatic network configuration is enabled, the command only allows you to enable or disable the SP network interface.

- If the SP automatic network configuration was previously enabled, disabling the SP network interface results in the assigned address resource being released and returned to the subnet.
- If you disable the SP network interface and then reenable it, the SP might be reconfigured with a different address.

If the SP automatic network configuration is disabled (the default), the following scenarios and considerations apply:

• If the SP has never been configured, SP IPv4 network configuration defaults to using IPv4 DHCP, and IPv6 is disabled.

A node that joins the cluster also uses IPv4 DHCP for its SP network configuration by default.

• The system service-processor network modify command enables you to configure a node's SP IP address.

A warning message appears when you attempt to manually configure the SP network with addresses that are allocated to a subnet. Ignoring the warning and proceeding with the manual address assignment might result in a scenario with duplicate addresses.

If the SP automatic network configuration is disabled after having been enabled previously, the following scenarios and considerations apply:

- If the SP automatic network configuration has the IPv4 address family disabled, the SP IPv4 network defaults to using DHCP, and the system service-processor network modify command enables you to modify the SP IPv4 configuration for individual nodes.
- If the SP automatic network configuration has the IPv6 address family disabled, the SP IPv6 network is also disabled, and the system service-processor network modify command enables you to enable and modify the SP IPv6 configuration for individual nodes.

Enable ONTAP SP/BMC automatic network configuration

Enabling the SP to use automatic network configuration is preferred over manually configuring the SP network. Because the SP automatic network configuration is cluster wide, you do not need to manually manage the SP network for individual nodes.



This task applies to both the SP and the BMC.

• The subnet you want to use for the SP automatic network configuration must already be defined in the cluster and must have no resource conflicts with the SP network interface.

The network subnet show command displays subnet information for the cluster.

Learn more about network subnet show in the ONTAP command reference.

The parameter that forces subnet association (the -force-update-lif-associations parameter of the network subnet commands) is supported only on network LIFs and not on the SP network interface.

• If you want to use IPv6 connections for the SP, IPv6 must already be configured and enabled for ONTAP.

The network options ipv6 show command displays the current state of IPv6 settings for ONTAP. Learn more about network options ipv6 show in the ONTAP command reference.

Steps

- 1. Specify the IPv4 or IPv6 address family and name for the subnet that you want the SP to use by using the system service-processor network auto-configuration enable command.
- 2. Display the SP automatic network configuration by using the system service-processor network auto-configuration show command.
- 3. If you subsequently want to disable or reenable the SP IPv4 or IPv6 network interface for all nodes that are in quorum, use the system service-processor network modify command with the -address -family [IPv4|IPv6] and -enable [true|false] parameters.

When the SP automatic network configuration is enabled, you cannot modify the SP IP address for a node

that is in quorum. You can only enable or disable the SP IPv4 or IPv6 network interface.

If a node is out of quorum, you can modify the node's SP network configuration, including the SP IP address, by running system service-processor network modify from the node and confirming that you want to override the SP automatic network configuration for the node. However, when the node joins the quorum, the SP automatic reconfiguration takes place for the node based on the specified subnet.

Configure the ONTAP SP/BMC network manually

If you do not have automatic network configuration set up for the SP, you must manually configure a node's SP network for the SP to be accessible by using an IP address.

Before you begin

If you want to use IPv6 connections for the SP, IPv6 must already be configured and enabled for ONTAP. The network options ipv6 commands manage IPv6 settings for ONTAP. Learn more about network options ipv6 in the ONTAP command reference.



This task applies to both the SP and the BMC.

You can configure the SP to use IPv4, IPv6, or both. The SP IPv4 configuration supports static and DHCP addressing, and the SP IPv6 configuration supports static addressing only.

If the SP automatic network configuration has been set up, you do not need to manually configure the SP network for individual nodes, and the system service-processor network modify command allows you to only enable or disable the SP network interface.

Steps

- 1. Configure the SP network for a node by using the system service-processor network modify command.
 - The -address-family parameter specifies whether the IPv4 or IPv6 configuration of the SP is to be modified.
 - ° The -enable parameter enables the network interface of the specified IP address family.
 - The -dhcp parameter specifies whether to use the network configuration from the DHCP server or the network address that you provide.

You can enable DHCP (by setting -dhcp to v4) only if you are using IPv4. You cannot enable DHCP for IPv6 configurations.

• The -ip-address parameter specifies the public IP address for the SP.

A warning message appears when you attempt to manually configure the SP network with addresses that are allocated to a subnet. Ignoring the warning and proceeding with the manual address assignment might result in a duplicate address assignment.

- The -netmask parameter specifies the netmask for the SP (if using IPv4.)
- The -prefix-length parameter specifies the network prefix-length of the subnet mask for the SP (if using IPv6.)
- The -gateway parameter specifies the gateway IP address for the SP.
- 2. Configure the SP network for the remaining nodes in the cluster by repeating the step 1.

3. Display the SP network configuration and verify the SP setup status by using the system serviceprocessor network show command with the -instance or -field setup-status parameters.

The SP setup status for a node can be one of the following:

- ° not-setup Not configured
- ° succeeded Configuration succeeded
- ° in-progress Configuration in progress
- ° failed Configuration failed

Example of configuring the SP network

The following example configures the SP of a node to use IPv4, enables the SP, and displays the SP network configuration to verify the settings:

```
cluster1::> system service-processor network modify -node local
-address-family IPv4 -enable true -ip-address 192.168.123.98
-netmask 255.255.255.0 -gateway 192.168.123.1
cluster1::> system service-processor network show -instance -node local
                               Node: node1
                       Address Type: IPv4
                  Interface Enabled: true
                     Type of Device: SP
                             Status: online
                        Link Status: up
                        DHCP Status: none
                         IP Address: 192.168.123.98
                        MAC Address: ab:cd:ef:fe:ed:02
                            Netmask: 255.255.255.0
       Prefix Length of Subnet Mask: -
         Router Assigned IP Address: -
              Link Local IP Address: -
                 Gateway IP Address: 192.168.123.1
                  Time Last Updated: Thu Apr 10 17:02:13 UTC 2014
                        Subnet Name: -
Enable IPv6 Router Assigned Address: -
            SP Network Setup Status: succeeded
    SP Network Setup Failure Reason: -
1 entries were displayed.
cluster1::>
```

Modify the ONTAP Service Processor API configuration

The SP API is a secure network API that enables ONTAP to communicate with the SP over the network. You can change the port used by the SP API service, renew the certificates the service uses for internal communication, or disable the service entirely. You need to modify the configuration only in rare situations.

About this task

• The SP API service uses port 50000 by default.

You can change the port value if, for example, you are in a network setting where port 50000 is used for communication by another networking application, or you want to differentiate between traffic from other applications and traffic generated by the SP API service.

• The SSL and SSH certificates used by the SP API service are internal to the cluster and not distributed externally.

In the unlikely event that the certificates are compromised, you can renew them.

• The SP API service is enabled by default.

You only need to disable the SP API service in rare situations, such as in a private LAN where the SP is not configured or used and you want to disable the service.

If the SP API service is disabled, the API does not accept any incoming connections. In addition, functionality such as network-based SP firmware updates and network-based SP "down system" log collection becomes unavailable. The system switches to using the serial interface.

Steps

- 1. Switch to the advanced privilege level by using the set -privilege advanced command.
- 2. Modify the SP API service configuration:

If you want to	Use the following command
Change the port used by the SP API service	system service-processor api-service modify with the -port {4915265535} parameter

If you want to	Use the following command
Renew the SSL and SSH certificates used by the SP API service for internal communication	• For ONTAP 9.5 or later use system service- processor api-service renew- internal-certificate
	 For ONTAP 9.4 and earlier use
	 system service-processor api- service renew-certificates
	If no parameter is specified, only the host certificates (including the client and server certificates) are renewed.
	If the -renew-all true parameter is specified, both the host certificates and the root CA certificate are renewed.
comm	
Disable or reenable the SP API service	system service-processor api-service modify with the -is-enabled {true false} parameter

3. Display the SP API service configuration by using the system service-processor api-service show command.

Manage nodes remotely using the SP/BMC

Manage an ONTAP node remotely using the SP/BMC

You can manage a node remotely using an onboard controller, called a Service Processor (SP) or Baseboard Management Controller (BMC). This remote management controller is included in all current platform models. The controller stays operational regardless of the operating state of the node.

For a full breakdown of platform SP and BMC support, refer to the Support Matrix on the NetApp Support site.

Remote node management with the ONTAP Service Processor

The Service Processor (SP) is a remote management device that enables you to access, monitor, and troubleshoot a node remotely.

The key capabilities of the SP include the following:

• The SP enables you to access a node remotely to diagnose, shut down, power-cycle, or reboot the node, regardless of the state of the node controller.

The SP is powered by a standby voltage, which is available as long as the node has input power from at least one of its power supplies.

You can log in to the SP by using a Secure Shell client application from an administration host. You can then use the SP CLI to monitor and troubleshoot the node remotely. In addition, you can use the SP to access the serial console and run ONTAP commands remotely.

You can access the SP from the serial console or access the serial console from the SP. The SP enables you to open both an SP CLI session and a separate console session simultaneously.

For instance, when a temperature sensor becomes critically high or low, ONTAP triggers the SP to shut down the motherboard gracefully. The serial console becomes unresponsive, but you can still press Ctrl-G on the console to access the SP CLI. You can then use the system power on or system power cycle command from the SP to power on or power-cycle the node.

• The SP monitors environmental sensors and logs events to help you take timely and effective service actions.

The SP monitors environmental sensors such as the node temperatures, voltages, currents, and fan speeds. When an environmental sensor has reached an abnormal condition, the SP logs the abnormal readings, notifies ONTAP of the issue, and sends alerts and "down system" notifications as necessary through an AutoSupport message, regardless of whether the node can send AutoSupport messages.

The SP also logs events such as boot progress, Field Replaceable Unit (FRU) changes, events generated by ONTAP, and SP command history. You can manually invoke an AutoSupport message to include the SP log files that are collected from a specified node.

Other than generating these messages on behalf of a node that is down and attaching additional diagnostic information to AutoSupport messages, the SP has no effect on the AutoSupport functionality. The AutoSupport configuration settings and message content behavior are inherited from ONTAP.



The SP does not rely on the -transport parameter setting of the system node autosupport modify command to send notifications. The SP only uses the Simple Mail Transport Protocol (SMTP) and requires its host's AutoSupport configuration to include mail host information.

If SNMP is enabled, the SP generates SNMP traps to configured trap hosts for all "down system" events.

• The SP has a nonvolatile memory buffer that stores up to 4,000 events in a system event log (SEL) to help you diagnose issues.

The SEL stores each audit log entry as an audit event. It is stored in onboard flash memory on the SP. The event list from the SEL is automatically sent by the SP to specified recipients through an AutoSupport message.

The SEL contains the following information:

- Hardware events detected by the SP—for example, sensor status about power supplies, voltage, or other components
- Errors detected by the SP—for example, a communication error, a fan failure, or a memory or CPU error
- Critical software events sent to the SP by the node—for example, a panic, a communication failure, a boot failure, or a user-triggered "down system" as a result of issuing the SP system reset or system power cycle command
- The SP monitors the serial console regardless of whether administrators are logged in or connected to the console.

When messages are sent to the console, the SP stores them in the console log. The console log persists as long as the SP has power from either of the node power supplies. Because the SP operates with standby power, it remains available even when the node is power-cycled or turned off.

- Hardware-assisted takeover is available if the SP is configured.
- The SP API service enables ONTAP to communicate with the SP over the network.

The service enhances ONTAP management of the SP by supporting network-based functionality such as using the network interface for the SP firmware update, enabling a node to access another node's SP functionality or system console, and uploading the SP log from another node.

You can modify the configuration of the SP API service by changing the port the service uses, renewing the SSL and SSH certificates that are used by the service for internal communication, or disabling the service entirely.

The following diagram illustrates access to ONTAP and the SP of a node. The SP interface is accessed through the Ethernet port (indicated by a wrench icon on the rear of the chassis):



Use the ONTAP Baseboard Management Controller to remotely manage a node

On certain hardware platforms, software is customized to support a new onboard controller in the Baseboard Management Controller (BMC). The BMC has command-line interface (CLI) commands you can use to manage the device remotely.

The BMC works similarly to the Service Processor (SP) and uses many of the same commands. The BMC allows you to do the following:

- Configure the BMC network settings.
- Access a node remotely and perform node management tasks such as diagnose, shut down, power-cycle, or reboot the node.

There are some differences between the SP and BMC:

- The BMC completely controls the environmental monitoring of power supply elements, cooling elements, temperature sensors, voltage sensors, and current sensors. The BMC reports sensor information to ONTAP through IPMI.
- Some of the high-availability (HA) and storage commands are different.
- The BMC does not send AutoSupport messages.

Automatic firmware updates are also available when running ONTAP with the following requirements:

• BMC firmware revision 1.15 or later must be installed.

()

A manual update is required to upgrade BMC firmware from 1.12 to 1.15 or later.

• BMC automatically reboots after a firmware update is completed.



Node operations are not impacted during a BMC reboot.

Manage ONTAP SP/BMC firmware updates

ONTAP includes an SP firmware image that is called the *baseline image*. If a new version of the SP firmware becomes subsequently available, you have the option to download it and update the SP firmware to the downloaded version without upgrading the ONTAP version.



This topic applies to both the SP and the BMC.

ONTAP offers the following methods for managing SP firmware updates:

- The SP automatic update functionality is enabled by default, allowing the SP firmware to be automatically updated in the following scenarios:
 - $\circ\,$ When you upgrade to a new version of ONTAP

The ONTAP upgrade process automatically includes the SP firmware update, provided that the SP firmware version bundled with ONTAP is newer than the SP version running on the node.



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 fail, see the Knowledge Base article xref:./system-admin/ Health Monitor SPAutoUpgradeFailedMajorAlert SP upgrade fails - AutoSupport Message.

- When you download a version of the SP firmware from the NetApp Support Site and the downloaded version is newer than the one that the SP is currently running
- · When you downgrade or revert to an earlier version of ONTAP

The SP firmware is automatically updated to the newest compatible version that is supported by the ONTAP version you reverted or downgraded to. A manual SP firmware update is not required.

You have the option to disable the SP automatic update functionality by using the system serviceprocessor image modify command. However, it is recommended that you leave the functionality enabled. Disabling the functionality can result in suboptimal or nonqualified combinations between the ONTAP image and the SP firmware image.

• ONTAP enables you to trigger an SP update manually and specify how the update should take place by using the system service-processor image update command.

You can specify the following options:

• The SP firmware package to use (-package)

You can update the SP firmware to a downloaded package by specifying the package file name. The advance system image package show command displays all package files (including the files for the SP firmware package) that are available on a node.

• Whether to use the baseline SP firmware package for the SP update (-baseline)

You can update the SP firmware to the baseline version that is bundled with the currently running version of ONTAP.



If you use some of the more advanced update options or parameters, the BMC's configuration settings may be temporarily cleared. After reboot, it can take up to 10 minutes for ONTAP to restore the BMC configuration.

• ONTAP enables you to display the status for the latest SP firmware update triggered from ONTAP by using the system service-processor image update-progress show command.

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 automatically or manually triggered.

Related information

NetApp Downloads: System Firmware and Diagnostics

ONTAP SP/BMC and network interface use for firmware updates

An SP firmware update that is triggered from ONTAP with the SP running version 1.5, 2.5, 3.1, or later supports using an IP-based file transfer mechanism over the SP network interface.



This topic applies to both the SP and the BMC.

An SP firmware update over the network interface is faster than an update over the serial interface. It reduces the maintenance window during which the SP firmware is being updated, and it is also nondisruptive to ONTAP operation. The SP versions that support this capability are included with ONTAP. They are also available on the NetApp Support Site and can be installed on controllers that are running a compatible version of ONTAP.

When you are running SP version 1.5, 2.5, 3.1, or later, the following firmware upgrade behaviors apply:

- An SP firmware update that is *automatically* triggered by ONTAP defaults to using the network interface for the update; however, the SP automatic update switches to using the serial interface for the firmware update if one of the following conditions occurs:
 - The SP network interface is not configured or not available.
 - The IP-based file transfer fails.
 - The SP API service is disabled.

Regardless of the SP version you are running, an SP firmware update triggered from the SP CLI always uses the SP network interface for the update.

Related information

NetApp Downloads: System Firmware and Diagnostics

Access the ONTAP Service Processor using a cluster user account

When you try to access the SP, you are prompted for credential. Cluster user accounts that are created with the service-processor application type have access to the SP CLI on any node of the cluster. SP user accounts are managed from ONTAP and authenticated by password. Beginning with ONTAP 9.9.1, SP user accounts must have the admin role.

User accounts for accessing the SP are managed from ONTAP instead of the SP CLI. A cluster user account can access the SP if it is created with the -application parameter of the security login create command set to service-processor and the -authmethod parameter set to password. The SP supports only password authentication.

You must specify the -role parameter when creating an SP user account.

- In ONTAP 9.9.1 and later releases, you must specify admin for the -role parameter, and any modifications to an account require the admin role. Other roles are no longer permitted for security reasons.
 - If you are upgrading to ONTAP 9.9.1 or later releases, see Change in user accounts that can access the Service Processor.
 - If you are reverting to ONTAP 9.8 or earlier releases, see Verify user accounts that can access the Service Processor.
- In ONTAP 9.8 and earlier releases, any role can access the SP, but admin is recommended.

By default, the cluster user account named "admin" includes the service-processor application type and has access to the SP.

ONTAP prevents you from creating user accounts with names that are reserved for the system (such as "root" and "naroot"). You cannot use a system-reserved name to access the cluster or the SP.

You can display current SP user accounts by using the -application service-processor parameter of the security login show command.

Learn more about security login show in the ONTAP command reference.

Access the ONTAP SP/BMC of a node from an administration host

You can log in to the SP of a node from an administration host to perform node management tasks remotely.

Before you begin

The following conditions must be met:

- The administration host you use to access the SP must support SSHv2.
- · Your user account must already be set up for accessing the SP.

To access the SP, your user account must have been created with the *-application* parameter of the security login create command set to service-processor and the *-authmethod* parameter set to password.



This task applies to both the SP and the BMC.

If the SP is configured to use an IPv4 or IPv6 address, and if five SSH login attempts from a host fail consecutively within 10 minutes, the SP rejects SSH login requests and suspends the communication with the IP address of the host for 15 minutes. The communication resumes after 15 minutes, and you can try to log in to the SP again.

ONTAP prevents you from creating or using system-reserved names (such as "root" and "naroot") to access the cluster or the SP.

Steps

1. From the administration host, log in to the SP:

ssh username@SP_IP_address

2. When you are prompted, enter the password for username.

The SP prompt appears, indicating that you have access to the SP CLI.

Examples of SP access from an administration host

The following example shows how to log in to the SP with a user account joe, which has been set up to access the SP.

```
[admin_host]$ ssh joe@192.168.123.98
joe@192.168.123.98's password:
SP>
```

The following examples show how to use the IPv6 global address or IPv6 router-advertised address to log in to the SP on a node that has SSH set up for IPv6 and the SP configured for IPv6.

```
[admin_host]$ ssh joe@fd22:8b1e:b255:202::1234
joe@fd22:8b1e:b255:202::1234's password:
SP>
```

```
[admin_host]$ ssh joe@fd22:8b1e:b255:202:2a0:98ff:fe01:7d5b
joe@fd22:8b1e:b255:202:2a0:98ff:fe01:7d5b's password:
SP>
```

Access the ONTAP SP/BMC of a node from the system console

You can access the SP from the system console (also called *serial console*) to perform monitoring or troubleshooting tasks.

About this task

This task applies to both the SP and the BMC.

Steps

- 1. Access the SP CLI from the system console by pressing Ctrl-G at the prompt.
- 2. Log in to the SP CLI when you are prompted.

The SP prompt appears, indicating that you have access to the SP CLI.

3. Exit the SP CLI and return to the system console by pressing Ctrl-D, and then press Enter.

Example of accessing the SP CLI from the system console

The following example shows the result of pressing Ctrl-G from the system console to access the SP CLI. The help system power command is entered at the SP prompt, followed by pressing Ctrl-D and then Enter to return to the system console.

cluster1::>

(Press Ctrl-G to access the SP CLI.)

```
Switching console to Service Processor
Service Processor Login:
Password:
SP>
SP> help system power
system power cycle - power the system off, then on
system power off - power the system off
system power on - power the system on
system power status - print system power status
SP>
```

(Press Ctrl-D and then Enter to return to the system console.)

cluster1::>

Learn how the ONTAP SP CLI, SP console, and system console sessions are related

You can open an SP CLI session to manage a node remotely and open a separate SP console session to access the console of the node. The SP console session mirrors output displayed in a concurrent system console session. The SP and the system console have independent shell environments with independent login authentication.

Understanding how the SP CLI, SP console, and system console sessions are related helps you manage a node remotely. The following describes the relationship among the sessions:

• Only one administrator can log in to the SP CLI session at a time; however, the SP enables you to open both an SP CLI session and a separate SP console session simultaneously.

The SP CLI is indicated with the SP prompt (SP>). From an SP CLI session, you can use the SP system

console command to initiate an SP console session. At the same time, you can start a separate SP CLI session through SSH. If you press Ctrl-D to exit from the SP console session, you automatically return to the SP CLI session. If an SP CLI session already exists, a message asks you whether to terminate the existing SP CLI session. If you enter "y", the existing SP CLI session is terminated, enabling you to return from the SP console to the SP CLI. This action is recorded in the SP event log.

In an ONTAP CLI session that is connected through SSH, you can switch to the system console of a node by running the ONTAP system node run-console command from another node.

• For security reasons, the SP CLI session and the system console session have independent login authentication.

When you initiate an SP console session from the SP CLI (by using the SP system console command), you are prompted for the system console credential. When you access the SP CLI from a system console session (by pressing Ctrl-G), you are prompted for the SP CLI credential.

• The SP console session and the system console session have independent shell environments.

The SP console session mirrors output that is displayed in a concurrent system console session. However, the concurrent system console session does not mirror the SP console session.

The SP console session does not mirror output of concurrent SSH sessions.

Add the administration host IP address to access the ONTAP Service Processor

By default, the SP accepts SSH connection requests from administration hosts of any IP addresses. You can configure the SP to accept SSH connection requests from only the administration hosts that have the IP addresses you specify. The changes you make apply to SSH access to the SP of any nodes in the cluster.

Steps

- 1. Grant SP access to only the IP addresses you specify by using the system service-processor ssh add-allowed-addresses command with the -allowed-addresses parameter.
 - The value of the -allowed-addresses parameter must be specified in the format of address /netmask, and multiple address/netmask pairs must be separated by commas, for example, 10.98.150.10/24, fd20:8ble:b255:c09b::/64.

Setting the -allowed-addresses parameter to 0.0.0/0, ::/0 enables all IP addresses to access the SP (the default).

- When you change the default by limiting SP access to only the IP addresses you specify, ONTAP prompts you to confirm that you want the specified IP addresses to replace the "allow all" default setting (0.0.0.0/0, ::/0).
- $^\circ$ The system service-processor ssh show command displays the IP addresses that can access the SP.
- 2. If you want to block a specified IP address from accessing the SP, use the system service-processor ssh remove-allowed-addresses command with the -allowed-addresses parameter.

If you block all IP addresses from accessing the SP, the SP becomes inaccessible from any administration hosts.

Examples of managing the IP addresses that can access the SP

The following examples show the default setting for SSH access to the SP, change the default by limiting SP access to only the specified IP addresses, remove the specified IP addresses from the access list, and then restore SP access for all IP addresses:

```
cluster1::> system service-processor ssh show
 Allowed Addresses: 0.0.0.0/0, ::/0
cluster1::> system service-processor ssh add-allowed-addresses -allowed
-addresses 192.168.1.202/24, 192.168.10.201/24
Warning: The default "allow all" setting (0.0.0.0/0, ::/0) will be
replaced
        with your changes. Do you want to continue? {y|n}: y
cluster1::> system service-processor ssh show
 Allowed Addresses: 192.168.1.202/24, 192.168.10.201/24
cluster1::> system service-processor ssh remove-allowed-addresses -allowed
-addresses 192.168.1.202/24, 192.168.10.201/24
Warning: If all IP addresses are removed from the allowed address list,
all IP
         addresses will be denied access. To restore the "allow all"
default,
         use the "system service-processor ssh add-allowed-addresses
         -allowed-addresses 0.0.0/0, ::/0" command. Do you want to
continue?
          {y|n}: y
cluster1::> system service-processor ssh show
 Allowed Addresses: -
cluster1::> system service-processor ssh add-allowed-addresses -allowed
-addresses 0.0.0/0, ::/0
cluster1::> system service-processor ssh show
 Allowed Addresses: 0.0.0.0/0, ::/0
```

View the help information in the ONTAP SP/BMC CLI

The online help displays the SP/BMC CLI commands and options.

About this task

This task applies to both the SP and the BMC.

Steps

1. To display help information for the SP/BMC commands, enter the following:

To access SP help	To access BMC help
Type help at the SP prompt.	Type system at the BMC prompt.

The following example shows the SP CLI online help.

SP> help date - print date and time exit - exit from the SP command line interface events - print system events and event information help - print command help priv - show and set user mode sp - commands to control the SP system - commands to control the system version - print SP version

The following example shows the BMC CLI online help.

```
BMC> system
system acp - acp related commands
system battery - battery related commands
system console - connect to the system console
system core - dump the system core and reset
system cpld - cpld commands
system log - print system console logs
system power - commands controlling system power
system reset - reset the system using the selected firmware
system sensors - print environmental sensors status
system fru - fru related commands
system watchdog - system watchdog commands
```

2. To display help information for the option of an SP/BMC command, enter help before or after the SP/BMC command.

The following example shows the SP CLI online help for the SP events command.

```
SP> help events
events all - print all system events
events info - print system event log information
events newest - print newest system events
events oldest - print oldest system events
events search - search for and print system events
```

The following example shows the BMC CLI online help for the BMC system power command.

```
BMC> system power help
system power cycle - power the system off, then on
system power off - power the system off
system power on - power the system on
system power status - print system power status
BMC>
```

Learn more about the commands described in this procedure in the ONTAP command reference.

Related information

- events
- system power

ONTAP commands for remote node management

You can manage a node remotely by accessing its SP and running SP CLI commands to perform node-management tasks. For several commonly performed remote node-management tasks, you can also use ONTAP commands from another node in the cluster. Some SP commands are platform-specific and might not be available on your platform.

If you want to	Use this SP command	Use this BMC command	Or this ONTAP command
Display available SP commands or subcommands of a specified SP command	help[command]		
Display the current privilege level for the SP CLI	priv show		

If you want to	Use this SP command	Use this BMC command	Or this ONTAP command
Set the privilege level to access the specified mode for the SP CLI	priv set{admin advanced diag}		
Display system date and time	date		date
Display events that are logged by the SP	events{all info newest number oldest number search keyword}		
Display SP status and network configuration information	sp status $[-v -d]$ The $-v$ option displays SP statistics in verbose form. The $-d$ option adds the SP debug log to the display.	bmc status $[-v -d]$ The $-v$ option displays SP statistics in verbose form. The $-d$ option adds the SP debug log to the display.	system service- processor show
Display the length of time the SP has been up and the average number of jobs in the run queue over the last 1, 5, and 15 minutes	sp uptime	bmc uptime	
Display system console logs	system log		
Display the SP log archives or the files in an archive	<pre>sp log history show [-archive {latest all archive-name}][-dump {all file- name}]</pre>	<pre>bmc log history show[-archive {latest all archive-name}][-dump {all file-name}]</pre>	
Display the power status for the controller of a node	system power status		system node power show
Display battery information	system battery show		
Display ACP information or the status for expander sensors	system acp[show sensors show]		

If you want to	Use this SP command	Use this BMC command	Or this ONTAP command
List all system FRUs and their IDs	system fru list		
Display product information for the specified FRU	system fru show fru_id		
Display the FRU data history log	system fru log show (advanced privilege level)		
Display the status for the environmental sensors, including their states and current values	system sensors Or system sensors show		system node environment sensors show
Display the status and details for the specified sensor	<pre>system sensors get sensor_name You can obtain sensor_name by using the system sensors or the system sensors show command.</pre>		
Display the SP firmware version information	version		system service- processor image show
Display the SP command history	sp log audit (advanced privilege level)	bmc log audit	
Display the SP debug information	sp log debug (advanced privilege level)	bmc log debug (advanced privilege level)	
Display the SP messages file	sp log messages (advanced privilege level)	bmc log messages (advanced privilege level)	

If you want to	Use this SP command	Use this BMC command	Or this ONTAP command
Display the settings for collecting system forensics on a watchdog reset event, display system forensics information collected during a watchdog reset event, or clear the collected system forensics information	system forensics [show log dump log clear]		
Log in to the system console	system console		system node run- console
	You should press Ctrl-D to	exit the system console ses	sion.
Turn the node on or off, or perform a power-cycle (turning the power off and then back on)	system power on		system node power on (advanced privilege level)
	system power off		
	system power cycle		
	The standby power stays on to keep the SP running without interruption. During the power-cycle, a brief pause occurs before power is turned back on.Image: Observed back of the power-cycle, a brief pause occurs before power is turned back on.Image: Observed back of the power-cycle, a brief pause occurs before power is turned back on.Image: Observed back of the power-cycle, a brief pause occurs before power is turned back on.Image: Observed back of the power-cycle, a brief pause occurs before power is turned back on.Image: Observed back of the power-cycle back of		
Create a core dump and reset the node	system core [-f] The -f option forces the creation of a core dump and the reset of the node.		system node coredump trigger (advanced privilege level)
	These commands have the same effect as pressing the Non-maskable Interrupt (NMI) button on a node, causing a dirty shutdown of the node and forcing a dump of the core files when halting the node. These commands are helpful when ONTAP on the node is hung or does not respond to commands such as system node shutdown. The generated core dump files are displayed in the output of the system node coredump show command. The SP stays operational as long as the input power to the node is not interrupted.		

If you want to	Use this SP command	Use this BMC command	Or this ONTAP command
Reboot the node with an optionally specified BIOS firmware image (primary, backup, or current) to recover from issues such as a corrupted image of the node's boot device	system reset {primary backup current}		<pre>system node reset with the -firmware {primary backup current} parameter(advanced privilege level) system node reset</pre>
	This operation	causes a dirty shutdown of	the node.
	If no BIOS firmware image The SP stays operational a interrupted.	is specified, the current ima s long as the input power to	ge is used for the reboot. the node is not
Display the status of battery firmware automatic update, or enable or disable battery firmware automatic update upon next SP boot	system battery auto_update[status enable disable] (advanced privilege level)		
Compare the current battery firmware image against a specified firmware image	<pre>system battery verify[image_URL] (advanced privilege level) If image_URL is not specified, the default battery firmware image is used for comparison.</pre>		
Update the battery firmware from the image at the specified location	system battery flash image_URL (advanced privilege level) You use this command if the automatic battery firmware upgrade process has failed for some reason.		
Update the SP firmware by using the image at the specified location	sp update image_URL image_URL must not exceed 200 characters .	bmc update image_URL image_URL must not exceed 200 characters.	system service- processor image update

If you want to	Use this SP command	Use this BMC command	Or this ONTAP command
Reboot the SP	sp reboot		system service- processor reboot-sp
Erase the NVRAM flash content	system nvram flash clear (advanced privilege level) This command cannot be initiated when the controller power is off (system power off).		
Exit the SP CLI	exit		

Learn more about the commands described in this procedure in the ONTAP command reference.

ONTAP node health monitoring using threshold-based SP sensor readings and status

Threshold-based sensors take periodic readings of a variety of system components. The SP compares the reading of a threshold-based sensor against its preset threshold limits that define a component's acceptable operating conditions.

Based on the sensor reading, the SP displays the sensor state to help you monitor the condition of the component.

Examples of threshold-based sensors include sensors for the system temperatures, voltages, currents, and fan speeds. The specific list of threshold-based sensors depends on the platform.

Threshold-based sensors have the following thresholds, displayed in the output of the SP system sensors command:

- Lower critical (LCR)
- Lower noncritical (LNC)
- Upper noncritical (UNC)
- Upper critical (UCR)

A sensor reading between LNC and LCR or between UNC and UCR means that the component is showing signs of a problem and a system failure might occur as a result. Therefore, you should plan for component service soon.

A sensor reading below LCR or above UCR means that the component is malfunctioning and a system failure is about to occur. Therefore, the component requires immediate attention.

The following diagram illustrates the severity ranges that are specified by the thresholds:


You can find the reading of a threshold-based sensor under the Current column in the system sensors command output. The system sensors get sensor_name command displays additional details for the specified sensor. As the reading of a threshold-based sensor crosses the noncritical and critical threshold ranges, the sensor reports a problem of increasing severity. When the reading exceeds a threshold limit, the sensor's status in the system sensors command output changes from ok to nc (noncritical) or cr (critical) depending on the exceeded threshold, and an event message is logged in the SEL event log.

Some threshold-based sensors do not have all four threshold levels. For those sensors, the missing thresholds show na as their limits in the system sensors command output, indicating that the particular sensor has no limit or severity concern for the given threshold and the SP does not monitor the sensor for that threshold.

Example of the system sensors command output

The following example shows some of the information displayed by the system sensors command in the SP CLI:

SP nodel> system sensors			
Sensor Name Current	Unit	Status LCR	LNC
	_+	_+	-+
++++++	· 		
CPU0_Temp_Margin -55.000	degrees C	ok na	na
-5.000 0.000 CPU1 Temp Margin -56.000	l degrees C	lok Ina	l na
-5.000 0.000	1 0.092000 0	, ,	1 110
In Flow Temp 32.000	degrees C	ok 0.000	10.000
42.000 52.000			
Out_Flow_Temp 38.000	degrees C	ok 0.000	10.000
59.000 68.000			
CPU1_Error 0x0	discrete	0x0180 na	na
na na			
CPU1_Therm_Trip 0x0	discrete	0x0180 na	na
na na			
CPU1_Hot 0x0	discrete	0x0180 na	na
na na	l dogrado C		
10_MIAI_TEMP 30.000	degrees c	OK 0.000	10.000
TO Mid2 Temp 30 000	l degrees C		
55.000 64.000	T degreeb e		1 10.000
CPU VTT 1.106	Volts	ok 1.028	1.048
1.154 1.174			
CPU0_VCC 1.154	Volts	ok 0.834	0.844
1.348 1.368			
3.3V 3.323	Volts	ok 3.053	3.116
3.466 3.546			
5V 5.002	Volts	ok 4.368	4.465
5.490 5.636			
STBY_1.8V 1.794	Volts	ok 1.678	1.707
1.892 1.911			

Example of the system sensors sensor_name command output for a threshold-based sensor

The following example shows the result of entering $system sensors get sensor_name in the SP CLI for the threshold-based sensor 5V:$

```
SP node1> system sensors get 5V
Locating sensor record...
             : 5V (0x13)
Sensor ID
Entity ID
                   : 7.97
Sensor Type (Analog) : Voltage
Sensor Reading : 5.002 (+/- 0) Volts
Status
                   : ok
Lower Non-Recoverable : na
Lower Critical : 4.246
Lower Non-Critical : 4.490
Upper Non-Critical
                   : 5.490
Upper Critical : 5.758
Upper Non-Recoverable : na
Assertion Events
                   :
Assertions Enabled : lnc- lcr- ucr+
Deassertions Enabled : lnc- lcr- ucr+
```

ONTAP SP sensor status values in the system sensor command output

Discrete sensors do not have thresholds. Their readings, displayed under the Current column in the SP CLI system sensors command output, do not carry actual meanings and thus are ignored by the SP. The Status column in the system sensors command output displays the status values of discrete sensors in hexadecimal format.

Examples of discrete sensors include sensors for the fan, power supply unit (PSU) fault, and system fault. The specific list of discrete sensors depends on the platform.

You can use the SP CLI system sensors get sensor_name command for help with interpreting the status values for most discrete sensors. The following examples show the results of entering system sensors get sensor_name for the discrete sensors CPU0_Error and IO_Slot1_Present:

```
SP nodel> system sensors get CPU0_Error
Locating sensor record...
Sensor ID : CPU0_Error (0x67)
Entity ID : 7.97
Sensor Type (Discrete): Temperature
States Asserted : Digital State
[State Deasserted]
```

```
SP nodel> system sensors get IO_Slot1_Present
Locating sensor record...
Sensor ID : IO_Slot1_Present (0x74)
Entity ID : 11.97
Sensor Type (Discrete): Add-in Card
States Asserted : Availability State
[Device Present]
```

Although the system sensors get sensor_name command displays the status information for most discrete sensors, it does not provide status information for the System_FW_Status, System_Watchdog, PSU1_Input_Type, and PSU2_Input_Type discrete sensors. You can use the following information to interpret these sensors' status values.

System_FW_Status

The System_FW_Status sensor's condition appears in the form of 0xAABB. You can combine the information of AA and BB to determine the condition of the sensor.

AA can have one of the following values:

Values	Condition of the sensor
01	System firmware error
02	System firmware hang
04	System firmware progress

BB can have one of the following values:

Values	Condition of the sensor
00	System software has properly shut down
01	Memory initialization in progress
02	NVMEM initialization in progress (when NVMEM is present)
04	Restoring memory controller hub (MCH) values (when NVMEM is present)
05	User has entered Setup
13	Booting the operating system or LOADER

Values	Condition of the sensor
1F	BIOS is starting up
20	LOADER is running
21	LOADER is programming the primary BIOS firmware. You must not power down the system.
22	LOADER is programming the alternate BIOS firmware. You must not power down the system.
2F	ONTAP is running
60	SP has powered off the system
61	SP has powered on the system
62	SP has reset the system
63	SP watchdog power cycle
64	SP watchdog cold reset

For instance, the System_FW_Status sensor status 0x042F means "system firmware progress (04), ONTAP is running (2F)."

System_Watchdog

The System_Watchdog sensor can have one of the following conditions:

• 0x0080

The state of this sensor has not changed

Values	Condition of the sensor
0x0081	Timer interrupt
0x0180	Timer expired
0x0280	Hard reset
0x0480	Power down
0x0880	Power cycle

For instance, the System_Watchdog sensor status 0x0880 means a watchdog timeout occurs and causes a system power cycle.

PSU1_Input_Type and PSU2_Input_Type

For direct current (DC) power supplies, the PSU1_Input_Type and PSU2_Input_Type sensors do not apply. For alternating current (AC) power supplies, the sensors' status can have one of the following values:

Values	Condition of the sensor
0x01 xx	220V PSU type
0x02 xx	110V PSU type

For instance, the PSU1_Input_Type sensor status 0x0280 means that the sensor reports that the PSU type is 110V.

ONTAP commands for Service Processor management

ONTAP provides commands for managing the SP, including the SP network configuration, SP firmware image, SSH access to the SP, and general SP administration.

Commands for managing the SP network configuration

If you want to	Run this ONTAP command
Enable the SP automatic network configuration for the SP to use the IPv4 or IPv6 address family of the specified subnet	system service-processor network auto- configuration enable
Disable the SP automatic network configuration for the IPv4 or IPv6 address family of the subnet specified for the SP	system service-processor network auto- configuration disable
Display the SP automatic network configuration	system service-processor network auto- configuration show

If you want to	Run this ONTAP command
Manually configure the SP network for a node, including the following:	system service-processor network modify
 The IP address family (IPv4 or IPv6) 	
 Whether the network interface of the specified IP address family should be enabled 	
 If you are using IPv4, whether to use the network configuration from the DHCP server or the network address that you specify 	
 The public IP address for the SP 	
 The netmask for the SP (if using IPv4) 	
 The network prefix-length of the subnet mask for the SP (if using IPv6) 	
 The gateway IP address for the SP 	
Display the SP network configuration, including the following:	system service-processor network show
lolowing.	Displaying complete SP network details requires the
 The configured address family (IPv4 or IPv6) and whether it is enabled 	-instance parameter.
 The remote management device type 	
 The current SP status and link status 	
 Network configuration, such as IP address, MAC address, netmask, prefix-length of subnet mask, router-assigned IP address, link local IP address, and gateway IP address 	
 The time the SP was last updated 	
 The name of the subnet used for SP automatic configuration 	
 Whether the IPv6 router-assigned IP address is enabled 	
 SP network setup status 	
 Reason for the SP network setup failure 	
Modify the SP API service configuration, including the following:	system service-processor api-service modify
 Changing the port used by the SP API service 	(advanced privilege level)
 Enabling or disabling the SP API service 	

If you want to	Run this ONTAP command
Display the SP API service configuration	system service-processor api-service show (advanced privilege level)
Renew the SSL and SSH certificates used by the SP API service for internal communication	• For ONTAP 9.5 or later: system service- processor api-service renew-internal- certificates
	• For ONTAP 9.4 or earlier: system service- processor api-service renew- certificates
	(advanced privilege level)

Commands for managing the SP firmware image

If you want to	Run this ONTAP command
 Display the details of the currently installed SP firmware image, including the following: The remote management device type The image (primary or backup) that the SP is booted from, its status, and firmware version Whether the firmware automatic update is enabled and the last update status 	system service-processor image show The -is-current parameter indicates the image (primary or backup) that the SP is currently booted from, not if the installed firmware version is most current.
Enable or disable the SP automatic firmware update	system service-processor image modify By default, the SP firmware is automatically updated with the update of ONTAP or when a new version of the SP firmware is manually downloaded. Disabling the automatic update is not recommended because doing so can result in suboptimal or nonqualified combinations between the ONTAP image and the SP firmware image.

If you want to	Run this ONTAP command	
Manually download an SP firmware image on a node	system node image get	
	(i) Before you run the system node image commands, you must set the privilege level to advanced (set -privilege advanced), entering y when prompted to continue.	
	The SP firmware image is packaged with ONTAP. You do not need to download the SP firmware manually, unless you want to use an SP firmware version that is different from the one packaged with ONTAP.	
Display the status for the latest SP firmware update triggered from ONTAP, including the following information:	system service-processor image update- progress show	
 The start and end time for the latest SP firmware update 		
 Whether an update is in progress and the percentage that is complete 		

Commands for managing SSH access to the SP

If you want to	Run this ONTAP command
Grant SP access to only the specified IP addresses	system service-processor ssh add- allowed-addresses
Block the specified IP addresses from accessing the SP	system service-processor ssh remove- allowed-addresses
Display the IP addresses that can access the SP	system service-processor ssh show

Commands for general SP administration

If you want to	Run this ONTAP command
Display general SP information, including the following:	system service-processor show Displaying complete SP information requires the -instance
 The remote management device type 	parameter.
The current SP status	
 Whether the SP network is configured 	
 Network information, such as the public IP address and the MAC address 	
 The SP firmware version and Intelligent Platform Management Interface (IPMI) version 	
 Whether the SP firmware automatic update is enabled 	
Reboot the SP on a node	system service-processor reboot-sp
Generate and send an AutoSupport message that includes the SP log files collected from a specified node	system node autosupport invoke-splog
Display the allocation map of the collected SP log files in the cluster, including the sequence numbers for the SP log files that reside in each collecting node	system service-processor log show- allocations

Related information

ONTAP command reference

ONTAP commands for BMC management

These ONTAP commands are supported on the Baseboard Management Controller (BMC).

The BMC uses some of the same commands as the Service Processor (SP). The following SP commands are supported on the BMC.

If you want to	Use this command
Display the BMC information	system service-processor show
Display/modify the BMC network configuration	system service-processor network show/modify
Reset the BMC	system service-processor reboot-sp

If you want to	Use this command
Display/modify the details of the currently installed BMC firmware image	system service-processor image show/modify
Update BMC firmware	system service-processor image update
Display the status for the latest BMC firmware update	system service-processor image update- progress show
Enable the automatic network configuration for the BMC to use an IPv4 or IPv6 address on the specified subnet	system service-processor network auto- configuration enable
Disable the automatic network configuration for an IPv4 or IPv6 address on the subnet specified for the BMC	system service-processor network auto- configuration disable
Display the BMC automatic network configuration	system service-processor network auto- configuration show

For commands that are not supported by the BMC firmware, the following error message is returned.

::> Error: Command not supported on this platform.

Related information

• system service-processor

Supported CLI commands for the ONTAP Baseboard Management Controller

You can log into the BMC using SSH. The following commands are supported from the BMC command line.

Command	Function
system	Display a list of all commands.
system console	Connect to the system's console. Use $\trl+D$ to exit the session.
system core	Dump the system core and reset.
system power cycle	Power the system off, then on.
system power off	Power the system off.

Command	Function
system power on	Power the system on.
system power status	Print system power status.
system reset	Reset the system.
system log	Print system console logs
system fru show [id]	Dump all/selected field replaceable unit (FRU) info.

Manage ONTAP cluster time (cluster administrators only)

Problems can occur when the cluster time is inaccurate. Although ONTAP enables you to manually set the time zone, date, and time on the cluster, you should configure the Network Time Protocol (NTP) servers to synchronize the cluster time.

Beginning with ONTAP 9.5, you can configure your NTP server with symmetric authentication.

NTP is always enabled. However, configuration is still required for the cluster to synchronize with an external time source. ONTAP enables you to manage the cluster's NTP configuration in the following ways:

- You can associate a maximum of 10 external NTP servers with the cluster (cluster time-service ntp server create).
 - For redundancy and quality of time service, you should associate at least three external NTP servers with the cluster.
 - You can specify an NTP server by using its IPv4 or IPv6 address or fully qualified host name.
 - $\,\circ\,$ You can manually specify the NTP version (v3 or v4) to use.

By default, ONTAP automatically selects the NTP version that is supported for a given external NTP server.

If the NTP version you specify is not supported for the NTP server, time exchange cannot take place.

- At the advanced privilege level, you can specify an external NTP server that is associated with the cluster to be the primary time source for correcting and adjusting the cluster time.
- You can display the NTP servers that are associated with the cluster (cluster time-service ntp server show).
- You can modify the cluster's NTP configuration (cluster time-service ntp server modify).
- You can disassociate the cluster from an external NTP server (cluster time-service ntp server delete).
- At the advanced privilege level, you can reset the configuration by clearing all external NTP servers' association with the cluster (cluster time-service ntp server reset).

A node that joins a cluster automatically adopts the NTP configuration of the cluster.

In addition to using NTP, ONTAP also enables you to manually manage the cluster time. This capability is helpful when you need to correct erroneous time (for example, a node's time has become significantly incorrect after a reboot). In that case, you can specify an approximate time for the cluster until NTP can synchronize with an external time server. The time you manually set takes effect across all nodes in the cluster.

You can manually manage the cluster time in the following ways:

- You can set or modify the time zone, date, and time on the cluster (cluster date modify).
- You can display the current time zone, date, and time settings of the cluster (cluster date show).



Job schedules do not adjust to manual cluster date and time changes. These jobs are scheduled to run based on the current cluster time when the job was created or when the job most recently ran. Therefore, if you manually change the cluster date or time, you must use the job show and job history show commands to verify that all scheduled jobs are queued and completed according to your requirements.

Commands for managing the cluster time

You use the cluster time-service ntp server commands to manage the NTP servers for the cluster. You use the cluster date commands to manage the cluster time manually.

Beginning with ONTAP 9.5, you can configure your NTP server with symmetric authentication.

The following commands enable you to manage the NTP servers for the cluster:
--

If you want to	Use this command
Associate the cluster with an external NTP server without symmetric authentication	cluster time-service ntp server create -server server_name
Associate the cluster with an external NTP server with symmetric authenticationAvailable in ONTAP 9.5 or later	cluster time-service ntp server create -server server_ip_address -key-id key_id
	(i) The key_id must refer to an existing shared key configured with '`cluster time-service ntp key'.
Enable symmetric authentication for an existing NTP serverAn existing NTP server can be modified to enable authentication by adding the required key-id. Available in ONTAP 9.5 or later	cluster time-service ntp server modify -server server_name -key-id key_id
Disable symmetric authentication	cluster time-service ntp server modify -server server_name -is-authentication -enabled false

If you want to	Use this command
Configure a shared NTP key	<pre>cluster time-service ntp key create -id shared_key_id -type shared_key_type -value shared_key_value</pre>
	Shared keys are referred to by an ID. The ID, its type, and value must be identical on both the node and the NTP server
Display information about the NTP servers that are associated with the cluster	cluster time-service ntp server show
Modify the configuration of an external NTP server that is associated with the cluster	cluster time-service ntp server modify
Dissociate an NTP server from the cluster	cluster time-service ntp server delete
Reset the configuration by clearing all external NTP servers' association with the cluster	cluster time-service ntp server reset
	This command requires the advanced privilege level.

The following commands enable you to manage the cluster time manually:

If you want to	Use this command
Set or modify the time zone, date, and time	cluster date modify
Display the time zone, date, and time settings for the cluster	cluster date show

Related information

- cluster date show
- cluster date modify
- cluster time-service ntp
- job show

Manage the banner and MOTD

Learn about the ONTAP login banner and message-of-the-day text

ONTAP enables you to configure a login banner or a message of the day (MOTD) to communicate administrative information to System Manager and CLI users of the cluster or storage virtual machine (SVM).

A banner is displayed in a console session (for cluster access only) or an SSH session (for cluster or SVM access) before a user is prompted for authentication such as a password. For example, you can use the banner to display a warning message such as the following to someone who attempts to log in to the system:

```
$ ssh admin@cluster1-01
This system is for authorized users only. Your IP Address has been logged.
Password:
```

An MOTD is displayed in a console session (for cluster access only) or an SSH session (for cluster or SVM access) after a user is authenticated but before the clustershell prompt appears. For example, you can use the MOTD to display a welcome or informational message such as the following that only authenticated users will see:

```
$ ssh admin@cluster1-01
Password:
Greetings. This system is running ONTAP 9.0.
Your user name is 'admin'. Your last login was Wed Apr 08 16:46:53 2015
from 10.72.137.28.
```

You can create or modify the content of the banner or MOTD by using the security login banner modify or security login motd modify command, respectively, in the following ways:

• You can use the CLI interactively or noninteractively to specify the text to use for the banner or MOTD.

The interactive mode, launched when the command is used without the *-message* or *-uri* parameter, enables you to use newlines (also known as end of lines) in the message.

The noninteractive mode, which uses the -message parameter to specify the message string, does not support newlines.

- You can upload content from an FTP or HTTP location to use for the banner or MOTD.
- You can configure the MOTD to display dynamic content.

Examples of what you can configure the MOTD to display dynamically include the following:

- · Cluster name, node name, or SVM name
- Cluster date and time
- Name of the user logging in
- Last login for the user on any node in the cluster
- Login device name or IP address
- Operating system name
- Software release version

• Effective cluster version string

The banner does not support dynamic content. Learn more about security login motd modify and the escape sequences that you can use to enable the MOTD to display dynamically generated content in the ONTAP command reference.

You can manage the banner and MOTD at the cluster or SVM level:

- The following facts apply to the banner:
 - The banner configured for the cluster is also used for all SVMs that do not have a banner message defined.
 - An SVM-level banner can be configured for each SVM.

If a cluster-level banner has been configured, it is overridden by the SVM-level banner for the given SVM.

- The following facts apply to the MOTD:
 - By default, the MOTD configured for the cluster is also enabled for all SVMs.
 - Additionally, an SVM-level MOTD can be configured for each SVM.

In this case, users logging in to the SVM will see two MOTDs, one defined at the cluster level and the other at the SVM level.

• The cluster-level MOTD can be enabled or disabled on a per-SVM basis by the cluster administrator.

If the cluster administrator disables the cluster-level MOTD for an SVM, a user logging in to the SVM does not see the cluster-level MOTD.

Create an ONTAP login banner

You can create a banner to display a message to someone who attempts to access the cluster or SVM. The banner is displayed in a console session (for cluster access only) or an SSH session (for cluster or SVM access) before a user is prompted for authentication.

Steps

1. Use the security login banner modify command to create a banner for the cluster or SVM:

If you want to	Then
Specify a message that is a single line	Use the -message " <text>" parameter to specify the text.</text>
Include newlines (also known as end of lines) in the message	Use the command without the -message or -uri parameter to launch the interactive mode for editing the banner.
Upload content from a location to use for the banner	Use the -uri parameter to specify the content's FTP or HTTP location.

The maximum size for a banner is 2,048 bytes, including newlines.

A banner created by using the -uri parameter is static. It is not automatically refreshed to reflect subsequent changes of the source content.

The banner created for the cluster is displayed also for all SVMs that do not have an existing banner. Any subsequently created banner for an SVM overrides the cluster-level banner for that SVM. Specifying the -message parameter with a hyphen within double quotes ("-") for the SVM resets the SVM to use the cluster-level banner.

2. Verify that the banner has been created by displaying it with the security login banner show command.

Specifying the -message parameter with an empty string ("") displays banners that have no content.

Specifying the -message parameter with "-" displays all (admin or data) SVMs that do not have a banner configured.

Examples of creating banners

The following example uses the noninteractive mode to create a banner for the "cluster1" cluster:

```
cluster1::> security login banner modify -message "Authorized users only!"
cluster1::>
```

The following example uses the interactive mode to create a banner for the svm1 SVM:

```
cluster1::> security login banner modify -vserver svm1
Enter the message of the day for Vserver "svm1".
Max size: 2048. Enter a blank line to terminate input. Press Ctrl-C to
abort.
                                                                   7
                  2
0
        1
                            3
                                      4
                                               5
                                                         6
8
12345678901234567890123456789012345678901234567890123456789012345678901234
567890
The svm1 SVM is reserved for authorized users only!
cluster1::>
```

The following example displays the banners that have been created:

```
cluster1::> security login banner show
Vserver: cluster1
Message
-----
Authorized users only!
Vserver: svm1
Message
-----
The svm1 SVM is reserved for authorized users only!
2 entries were displayed.
cluster1::>
```

Related information

- Managing the banner
- security login banner modify
- security login banner show

Manage the banner text displayed at the ONTAP cluster and SVM level

You can manage the banner at the cluster or SVM level. The banner configured for the cluster is also used for all SVMs that do not have a banner message defined. A subsequently created banner for an SVM overrides the cluster banner for that SVM.

Choices

• Manage the banner at the cluster level:

If you want to	Then
Create a banner to display for all login sessions	<pre>Set a cluster-level banner: security login banner modify -vserver <cluster_name> { [-message "text"] [-uri <ftp_or_http_addr>] }</ftp_or_http_addr></cluster_name></pre>
Remove the banner for all (cluster and SVM) logins	Set the banner to an empty string (``): security login banner modify -vserver * -message

If you want to	Then
Override a banner created by an SVM administrator	<pre>Modify the SVM banner message: security login banner modify -vserver <svm_name> { [-message "<text>"] [- uri <ftp_or_http_addr>] }</ftp_or_http_addr></text></svm_name></pre>

• Manage the banner at the SVM level:

Specifying -vserver <svm name> is not required in the SVM context.

If you want to	Then
Override the banner supplied by the cluster administrator with a different banner for the SVM	<pre>Create a banner for the SVM: security login banner modify -vserver <svm_name> { [-message "text"] [-uri <ftp_or_http_addr>] }</ftp_or_http_addr></svm_name></pre>
Suppress the banner supplied by the cluster administrator so that no banner is displayed for the SVM	Set the SVM banner to an empty string for the SVM: security login banner modify -vserver <svm_name> -message</svm_name>
Use the cluster-level banner when the SVM currently uses an SVM-level banner	Set the SVM banner to -: security login banner modify -vserver <svm_name> -message -</svm_name>

Related information

• security login banner modify

Create message-of-the-day text for ONTAP users

You can create a message of the day (MOTD) to communicate information to authenticated CLI users. The MOTD is displayed in a console session (for cluster access only) or an SSH session (for cluster or SVM access) after a user is authenticated but before the clustershell prompt appears.

Steps

1. Use the security login motd modify command to create an MOTD for the cluster or SVM:

If you want to	Then
Specify a message that is a single line	Use the -message "text" parameter to specify the text.

If you want to	Then
Include newlines (also known as end of lines)	Use the command without the -message or -uri parameter to launch the interactive mode for editing the MOTD.
Upload content from a location to use for the MOTD	Use the -uri parameter to specify the content's FTP or HTTP location.

The maximum size for an MOTD is 2,048 bytes, including newlines.

Security login motd modify describes the escape sequences that you can use to enable the MOTD to display dynamically generated content.

An MOTD created by using the -uri parameter is static. It is not automatically refreshed to reflect subsequent changes of the source content.

An MOTD created for the cluster is displayed also for all SVM logins by default, along with an SVM-level MOTD that you can create separately for a given SVM. Setting the <code>-is-cluster-message-enabled</code> parameter to <code>false</code> for an SVM prevents the cluster-level MOTD from being displayed for that SVM.

2. Verify that the MOTD has been created by displaying it with the security login motd show command.

Specifying the -message parameter with an empty string (``) displays MOTDs that are not configured or have no content.

Learn more about security login motd modify and parameters to enable the MOTD to display dynamically generated content in the ONTAP command reference.

Examples of creating MOTDs

The following example uses the noninteractive mode to create an MOTD for the cluster1 cluster:

```
cluster1::> security login motd modify -message "Greetings!"
```

The following example uses the interactive mode to create an MOTD for the svm1 SVM that uses escape sequences to display dynamically generated content:

The following example displays the MOTDs that have been created:

```
cluster1::> security login motd show
Vserver: cluster1
Is the Cluster MOTD Displayed?: true
Message
-----
Greetings!
Vserver: svm1
Is the Cluster MOTD Displayed?: true
Message
-----
Welcome to the \n SVM. Your user ID is '\N'. Your last successful login
was \L.
2 entries were displayed.
```

Learn more about security login motd show in the ONTAP command reference.

Manage the message-of-the-day text displayed at ONTAP cluster and SVM levels

You can manage the message of the day (MOTD) at the cluster or SVM level. By default, the MOTD configured for the cluster is also enabled for all SVMs. Additionally, an SVM-level MOTD can be configured for each SVM. The cluster-level MOTD can be enabled or disabled for each SVM by the cluster administrator.

Learn more about the escape sequences that can be used to dynamically generate content for the MOTD in the ONTAP command reference.

Choices

• Manage the MOTD at the cluster level:

If you want to	Then
Create an MOTD for all logins when there is no existing MOTD	<pre>Set a cluster-level MOTD: security login motd modify -vserver <cluster_name> { [-message "<text>"] [-uri <ftp_or_http_addr>] }</ftp_or_http_addr></text></cluster_name></pre>
Change the MOTD for all logins when no SVM-level MOTDs are configured	<pre>Modify the cluster-level MOTD: security login motd modify -vserver <cluster_name> { [-message "<text>"] } [-uri <ftp_or_http_addr>] }</ftp_or_http_addr></text></cluster_name></pre>
Remove the MOTD for all logins when no SVM-level MOTDs are configured	Set the cluster-level MOTD to an empty string (""): security login motd modify -vserver <cluster_name> -message ""</cluster_name>
Have every SVM display the cluster-level MOTD instead of using the SVM-level MOTD	<pre>Set a cluster-level MOTD, then set all SVM-level MOTDs to an empty string with the cluster-level MOTD enabled: 1. security login motd modify -vserver</pre>
Have an MOTD displayed for only selected SVMs, and use no cluster-level MOTD	<pre>Set the cluster-level MOTD to an empty string, then set SVM-level MOTDs for selected SVMs: 1. security login motd modify -vserver <cluster_name> -message "" 2. security login motd modify -vserver <svm_name> { [-message "<text>"] [-uri <ftp_or_http_addr>] } You can repeat this step for each SVM as needed.</ftp_or_http_addr></text></svm_name></cluster_name></pre>

If you want to	Then
Use the same SVM-level MOTD for all (data and admin) SVMs	Set the cluster and all SVMs to use the same MOTD:
	<pre>security login motd modify -vserver * { [-message "<text>"] [-uri <ftp_or_http_addr>] }</ftp_or_http_addr></text></pre>
	(i) If you use the interactive mode, the CLI prompts you to enter the MOTD individually for the cluster and each SVM. You can paste the same MOTD into each instance when you are prompted to.
Have a cluster-level MOTD optionally available to all SVMs, but do not want the MOTD displayed for cluster logins	Set a cluster-level MOTD, but disable its display for the cluster:
	<pre>security login motd modify -vserver <cluster_name> { [-message "<text>"] [-uri <ftp_or_http_addr>] } -is -cluster-message-enabled false</ftp_or_http_addr></text></cluster_name></pre>
Remove all MOTDs at the cluster and SVM levels when only some SVMs have both cluster-level and SVM-level MOTDs	Set the cluster and all SVMs to use an empty string for the MOTD: security login motd modify -vserver *
	-message ""
Modify the MOTD only for the SVMs that have a non-empty string, when other SVMs use an empty string, and when a different MOTD is used at the cluster level	Use extended queries to modify the MOTD selectively:
	<pre>security login motd modify { -vserver !"<cluster_name>" -message !"" } { [- message "<text>"] [-uri <ftp_or_http_addr>] }</ftp_or_http_addr></text></cluster_name></pre>
Display all MOTDs that contain specific text (for example, "January" followed by "2015") anywhere in a single or multiline message, even if the text is split across different lines	Use a query to display MOTDs:
	security login motd show -message *"January"*"2015"*
Interactively create an MOTD that includes multiple and consecutive newlines (also known as end of lines, or EOLs)	In the interactive mode, press the space bar followed by Enter to create a blank line without terminating the input for the MOTD.

• Manage the MOTD at the SVM level:

If you want to	Then
Use a different SVM-level MOTD, when the SVM already has an existing SVM-level MOTD	<pre>Modify the SVM-level MOTD: security login motd modify -vserver <svm_name> { [-message "<text>"] [- uri <ftp_or_http_addr>] }</ftp_or_http_addr></text></svm_name></pre>
Use only the cluster-level MOTD for the SVM, when the SVM already has an SVM-level MOTD	<pre>Set the SVM-level MOTD to an empty string, then have the cluster administrator enable the cluster- level MOTD for the SVM: 1. security login motd modify -vserver</pre>
Not have the SVM display any MOTD, when both the cluster-level and SVM-level MOTDs are currently displayed for the SVM	<pre>Set the SVM-level MOTD to an empty string, then have the cluster administrator disable the cluster- level MOTD for the SVM: 1. security login motd modify -vserver</pre>

Related information

- security login motd modify
- security login motd show

Manage ONTAP jobs and job schedules

Jobs are placed into a job queue and run in the background when resources are available. If a job is consuming too many cluster resources, you can stop it or pause it until there is less demand on the cluster. You can also monitor and restart jobs.

Job categories

There are three categories of jobs that you can manage: server-affiliated, cluster-affiliated, and private.

A job can be in any of the following categories:

· Server-Affiliated jobs

These jobs are queued by the management framework to a specific node to be run.

Cluster-Affiliated jobs

These jobs are queued by the management framework to any node in the cluster to be run.

Private jobs

These jobs are specific to a node and do not use the replicated database (RDB) or any other cluster mechanism. The commands that manage private jobs require the advanced privilege level or higher.

Commands for managing jobs

When you enter a command that invokes a job, typically, the command informs you that the job has been queued and then returns to the CLI command prompt. However, some commands instead report job progress and do not return to the CLI command prompt until the job has been completed. In these cases, you can press Ctrl-C to move the job to the background.

If you want to	Use this command
Display information about all jobs	job show
Display information about jobs on a per-node basis	job show bynode
Display information about cluster-affiliated jobs	job show-cluster
Display information about completed jobs	job show-completed
Display information about job history	job history show Up to 25,000 job records are stored for each node in the cluster. Consequently, attempting to display the full job history could take a long time. To avoid potentially long wait times, you should display jobs by node, storage virtual machine (SVM), or record ID.
Display the list of private jobs	job private show (advanced privilege level)
Display information about completed private jobs	job private show-completed (advanced privilege level)
Display information about the initialization state for job managers	job initstate show (advanced privilege level)
Monitor the progress of a job	job watch-progress

If you want to	Use this command
Monitor the progress of a private job	job private watch-progress (advanced privilege level)
Pause a job	job pause
Pause a private job	job private pause (advanced privilege level)
Resume a paused job	job resume
Resume a paused private job	job private resume (advanced privilege level)
Stop a job	job stop
Stop a private job	job private stop (advanced privilege level)
Delete a job	job delete
Delete a private job	job private delete (advanced privilege level)
Disassociate a cluster-affiliated job with an unavailable node that owns it, so that another node can take ownership of that job	job unclaim (advanced privilege level)



You can use the event log show command to determine the outcome of a completed job. Learn more about event log show in the ONTAP command reference.

Commands for managing job schedules

Many tasks—for instance, volume snapshots—can be configured to run on specified schedules.Schedules that run at specific times are called *cron* schedules (similar to UNIX cron schedules). Schedules that run at intervals are called *interval* schedules. You use the job schedule commands to manage job schedules.

Job schedules do not adjust to manual changes to the cluster date and time. These jobs are scheduled to run based on the current cluster time when the job was created or when the job most recently ran. Therefore, if you manually change the cluster date or time, you should use the job show and job history show commands to verify that all scheduled jobs are queued and completed according to your requirements.

If the cluster is part of a MetroCluster configuration, then the job schedules on both clusters must be identical. Therefore, if you create, modify, or delete a job schedule, you must perform the same operation on the remote cluster.

If you want to	Use this command
Display information about all schedules	job schedule show

If you want to	Use this command
Display the list of jobs by schedule	job schedule show-jobs
Display information about cron schedules	job schedule cron show
Display information about interval schedules	job schedule interval show
Create a cron schedule	job schedule cron create Beginning with ONTAP 9.10.1, you can include the SVM for your job schedule.
Create an interval schedule	job schedule interval create You must specify at least one of the following parameters: -days, -hours, -minutes, or -seconds.
Modify a cron schedule	job schedule cron modify
Modify an interval schedule	job schedule interval modify
Delete a schedule	job schedule delete
Delete a cron schedule	job schedule cron delete
Delete an interval schedule	job schedule interval delete

Related information

• job

Back up and restore cluster configurations (cluster administrators only)

Learn about ONTAP configuration backup files

Configuration backup files are archive files (.7z) that contain information for all configurable options that are necessary for the cluster, and the nodes within it, to operate properly.

These files store the local configuration of each node, plus the cluster-wide replicated configuration. You use configuration backup files to back up and restore the configuration of your cluster.

There are two types of configuration backup files:

Node configuration backup file

Each healthy node in the cluster includes a node configuration backup file, which contains all of the configuration information and metadata necessary for the node to operate healthy in the cluster.

Cluster configuration backup file

These files include an archive of all of the node configuration backup files in the cluster, plus the replicated cluster configuration information (the replicated database, or RDB file). Cluster configuration backup files enable you to restore the configuration of the entire cluster, or of any node in the cluster. The cluster configuration backup schedules create these files automatically and store them on several nodes in the cluster.



Configuration backup files contain configuration information only. They do not include any user data. For information about restoring user data, see Data Protection.

Learn about scheduling backups of ONTAP cluster and node configuration backup files

Three separate schedules automatically create cluster and node configuration backup files and replicate them among the nodes in the cluster.

The configuration backup files are automatically created according to the following schedules:

- Every 8 hours
- Daily
- Weekly

At each of these times, a node configuration backup file is created on each healthy node in the cluster. All of these node configuration backup files are then collected in a single cluster configuration backup file along with the replicated cluster configuration and saved on one or more nodes in the cluster.

ONTAP commands for managing configuration backup schedules

You can use the system configuration backup settings commands to manage configuration backup schedules.

These commands are available at the advanced privilege level.

If you want to	Use this command
Change the settings for a configuration backup schedule:	system configuration backup settings modify
 Specify a remote URL (HTTP, HTTPS, FTP, FTPS, or TFTP) where the configuration backup files will be uploaded in addition to the default locations in the cluster 	When you use HTTPS in the remote URL, use the -validate-certification option to enable or disable digital certificate validation. Certificate validation is disabled by default.
 Specify a user name to be used to log in to the remote URL 	The web server to which you are
 Set the number of backups to keep for each configuration backup schedule 	(i) uploading the configuration backup file must have PUT operations enabled for HTTP and POST operations enabled for HTTPS. For more information, see your web server's documentation.
Set the password to be used to log in to the remote URL	system configuration backup settings set-password
View the settings for the configuration backup schedule	system configuration backup settings show
	(i) You set the -instance parameter to view the user name and the number of backups to keep for each schedule.

ONTAP commands for managing node configuration backup files

You use the system configuration backup commands to manage cluster and node configuration backup files.

These commands are available at the advanced privilege level.

If you want to	Use this command
Create a new node or cluster configuration backup file	system configuration backup create
Copy a configuration backup file from a node to another node in the cluster	system configuration backup copy

If you want to	Use this command		
Upload a configuration backup file from a node in the cluster to a remote URL (FTP, HTTP, HTTPS, TFTP, or FTPS)	system configuration backup uploadWhen you use HTTPS in the remote URL, use the -validate-certification option to enable or disable digital certificate validation. Certificate validation is disabled by default.The web server to which you are uploading the configuration backup file must have PUT operations enabled for HTTP and POST operations enabled 		
	web server's documentation. Supported URL formats vary by ONTAP release. Learn more about system configuration commands in the `ONTAP command reference.		
Download a configuration backup file from a remote URL to a node in the cluster, and, if specified, validate the digital certificate	system configuration backup download When you use HTTPS in the remote URL, use the -validate-certification option to enable or disable digital certificate validation. Certificate validation is disabled by default.		
Rename a configuration backup file on a node in the cluster	system configuration backup rename		
View the node and cluster configuration backup files for one or more nodes in the cluster	system configuration backup show		
Delete a configuration backup file on a node	system configuration backup deleteThis command deletes the configuration backup file on the specified node only. If the configuration backup file also exists on other nodes in the cluster, it remains on those nodes.		

Locate the ONTAP node configuration backup file to restore a node

You use a configuration backup file located at a remote URL or on a node in the cluster to recover a node configuration.

About this task

You can use either a cluster or node configuration backup file to restore a node configuration.

Step

1. Make the configuration backup file available to the node for which you need to restore the configuration.

If the configuration backup file is located	Then	
At a remote URL	Use the system configuration backup download command at the advanced privilege level to download it to the recovering node.	
On a node in the cluster	a. Use the system configuration backup show command at the advanced privilege level to view the list of configuration backup files available in the cluster that contains the recovering node's configuration.	
	b. If the configuration backup file you identify does not exist on the recovering node, then use the system configuration backup copy command to copy it to the recovering node.	

If you previously re-created the cluster, you should choose a configuration backup file that was created after the cluster recreation. If you must use a configuration backup file that was created prior to the cluster recreation, then after recovering the node, you must re-create the cluster again.

Restore a node using the ONTAP node configuration backup file

You restore the node configuration using the configuration backup file that you identified and made available to the recovering node.

About this task

You should only perform this task to recover from a disaster that resulted in the loss of the node's local configuration files.

Steps

1. Change to the advanced privilege level:

set -privilege advanced

2. If the node is healthy, then at the advanced privilege level of a different node, use the cluster modify command with the -node and -eligibility parameters to mark it ineligible and isolate it from the cluster.

If the node is not healthy, then you should skip this step.

This example modifies node2 to be ineligible to participate in the cluster so that its configuration can be restored:

cluster1::*> cluster modify -node node2 -eligibility false

Learn more about cluster modify in the ONTAP command reference.

3. Use the system configuration recovery node restore command at the advanced privilege level to restore the node's configuration from a configuration backup file.

If the node lost its identity, including its name, then you should use the *-nodename-in-backup* parameter to specify the node name in the configuration backup file.

This example restores the node's configuration using one of the configuration backup files stored on the node:

```
cluster1::*> system configuration recovery node restore -backup
cluster1.8hour.2011-02-22.18_15_00.7z
Warning: This command overwrites local configuration files with
files contained in the specified backup file. Use this
command only to recover from a disaster that resulted
in the loss of the local configuration files.
The node will reboot after restoring the local configuration.
Do you want to continue? {y|n}: y
```

The configuration is restored, and the node reboots.

- 4. If you marked the node ineligible, then use the system configuration recovery cluster sync command to mark the node as eligible and synchronize it with the cluster.
- 5. If you are operating in a SAN environment, use the system node reboot command to reboot the node and reestablish SAN quorum.

After you finish

If you previously re-created the cluster, and if you are restoring the node configuration by using a configuration backup file that was created prior to that cluster re-creation, then you must re-create the cluster again.

Locate the ONTAP cluster configuration backup file to restore a cluster

You use the configuration from either a node in the cluster or a cluster configuration backup file to recover a cluster.

Steps

- 1. Choose a type of configuration to recover the cluster.
 - A node in the cluster

If the cluster consists of more than one node, and one of the nodes has a cluster configuration from when the cluster was in the desired configuration, then you can recover the cluster using the configuration stored on that node.

In most cases, the node containing the replication ring with the most recent transaction ID is the best node to use for restoring the cluster configuration. The cluster ring show command at the advanced privilege level enables you to view a list of the replicated rings available on each node in the cluster.

• A cluster configuration backup file

If you cannot identify a node with the correct cluster configuration, or if the cluster consists of a single node, then you can use a cluster configuration backup file to recover the cluster.

If you are recovering the cluster from a configuration backup file, any configuration changes made since the backup was taken will be lost. You must resolve any discrepancies between the configuration backup file and the present configuration after recovery. See Knowledge Base article ONTAP Configuration Backup Resolution Guide for troubleshooting guidance.

2. If you chose to use a cluster configuration backup file, then make the file available to the node you plan to use to recover the cluster.

If the configuration backup file is located	Then	
At a remote URL	Use the system configuration backup download command at the advanced privilege level to download it to the recovering node.	
On a node in the cluster	a. Use the system configuration backup show command at the advanced privilege level to find a cluster configuration backup file that was created when the cluster was in the desired configuration.	
	b. If the cluster configuration backup file is not located on the node you plan to use to recover the cluster, then use the system configuration backup copy command to copy it to the recovering node.	

Learn more about cluster ring show in the ONTAP command reference.

Restore a cluster using the ONTAP cluster configuration backup file

To restore a cluster configuration from an existing configuration after a cluster failure, you re-create the cluster using the cluster configuration that you chose and made available to the recovering node, and then rejoin each additional node to the new cluster.

About this task

You should only perform this task to recover from a disaster that resulted in the loss of the cluster's configuration.

If you are re-creating the cluster from a configuration backup file, you must contact technical support to resolve any discrepancies between the configuration backup file and the configuration present in the cluster.

If you are recovering the cluster from a configuration backup file, any configuration changes made since the backup was taken will be lost. You must resolve any discrepancies between the configuration backup file and the present configuration after recovery. See the Knowledge Base article ONTAP Configuration Backup Resolution Guide for troubleshooting guidance.

Steps

1. Disable storage failover for each HA pair:

storage failover modify -node node_name -enabled false

You only need to disable storage failover once for each HA pair. When you disable storage failover for a node, storage failover is also disabled on the node's partner.

2. Halt each node except for the recovering node:

```
system node halt -node node name -reason "text"
```

```
cluster1::*> system node halt -node node0 -reason "recovering cluster"
Warning: Are you sure you want to halt the node? {y|n}: y
```

3. Set the privilege level to advanced:

set -privilege advanced

 On the recovering node, use the system configuration recovery cluster recreate command to re-create the cluster.

This example re-creates the cluster using the configuration information stored on the recovering node:

```
cluster1::*> configuration recovery cluster recreate -from node
Warning: This command will destroy your existing cluster. It will
rebuild a new single-node cluster consisting of this node
and its current configuration. This feature should only be
used to recover from a disaster. Do not perform any other
recovery operations while this operation is in progress.
Do you want to continue? {y|n}: y
```

A new cluster is created on the recovering node.

If you are re-creating the cluster from a configuration backup file, verify that the cluster recovery is still in progress:

system configuration recovery cluster show

You do not need to verify the cluster recovery state if you are re-creating the cluster from a healthy node.

```
cluster1::*> system configuration recovery cluster show
Recovery Status: in-progress
Is Recovery Status Persisted: false
```

6. Boot each node that needs to be rejoined to the re-created cluster.

You must reboot the nodes one at a time.

- 7. For each node that needs to be joined to the re-created cluster, do the following:
 - a. From a healthy node on the re-created cluster, rejoin the target node:

system configuration recovery cluster rejoin -node node_name

This example rejoins the "node2" target node to the re-created cluster:

```
cluster1::*> system configuration recovery cluster rejoin -node node2
Warning: This command will rejoin node "node2" into the local
cluster, potentially overwriting critical cluster
configuration files. This command should only be used
to recover from a disaster. Do not perform any other
recovery operations while this operation is in progress.
This command will cause node "node2" to reboot.
Do you want to continue? {y|n}: y
```

The target node reboots and then joins the cluster.

b. Verify that the target node is healthy and has formed quorum with the rest of the nodes in the cluster:

cluster show -eligibility true

The target node must rejoin the re-created cluster before you can rejoin another node.

<pre>cluster1::*> cluster</pre>	show -eligibility true				
Node	Health	Eligibility	Epsilon		
node0	true	true	false		
nodel	true	true	false		
2 entries were displayed.					

8. If you re-created the cluster from a configuration backup file, set the recovery status to be complete:

system configuration recovery cluster modify -recovery-status complete

9. Return to the admin privilege level:

set -privilege admin

10. If the cluster consists of only two nodes, use the **cluster ha modify** command to reenable cluster HA.

11. Use the **storage failover modify** command to reenable storage failover for each HA pair.

After you finish

If the cluster has SnapMirror peer relationships, then you also need to re-create those relationships. For more information, see Data Protection.

Related information

• ONTAP command reference

Synchronize a node with the ONTAP cluster to ensure cluster-wide quorum

If cluster-wide quorum exists, but one or more nodes are out of sync with the cluster, then you must synchronize the node to restore the replicated database (RDB) on the node and bring it into quorum.

Step

1. From a healthy node, use the system configuration recovery cluster sync command at the advanced privilege level to synchronize the node that is out of sync with the cluster configuration.

This example synchronizes a node (node2) with the rest of the cluster:

```
cluster1::*> system configuration recovery cluster sync -node node2
Warning: This command will synchronize node "node2" with the cluster
    configuration, potentially overwriting critical cluster
    configuration files on the node. This feature should only be
    used to recover from a disaster. Do not perform any other
    recovery operations while this operation is in progress. This
    command will cause all the cluster applications on node
    "node2" to restart, interrupting administrative CLI and Web
    interface on that node.
Do you want to continue? {y|n}: y
All cluster applications on node "node2" will be restarted. Verify that
    the cluster applications go online.
```

Result

The RDB is replicated to the node, and the node becomes eligible to participate in the cluster.

Manage node core dumps for an ONTAP cluster (cluster administrators only)

When a node panics, a core dump occurs and the system creates a core dump file that technical support can use to troubleshoot the problem. You can configure or display core dump attributes. You can also save, display, segment, upload, or delete a core dump file.

You can manage core dumps in the following ways:

- · Configuring core dumps and displaying the configuration settings
- · Displaying basic information, the status, and attributes of core dumps

Core dump files and reports are stored in the /mroot/etc/crash/ directory of a node. You can display
the directory content by using the system node coredump commands or a web browser.

• Saving the core dump content and uploading the saved file to a specified location or to technical support

ONTAP prevents you from initiating the saving of a core dump file during a takeover, an aggregate relocation, or a giveback.

• Deleting core dump files that are no longer needed

Commands for managing core dumps

You use the system node coredump config commands to manage the configuration of core dumps, the system node coredump commands to manage the core dump files, and the system node coredump reports commands to manage application core reports.

Learn more about the commands described in this topic in the ONTAP command reference.

If you want to	Use this command		
Configure core dumps	system node coredump config modify		
Display the configuration settings for core dumps	system node coredump config show		
Display basic information about core dumps	system node coredump show		
Manually trigger a core dump when you reboot a node	system node reboot with both the -dump and -skip-lif-migration-before-reboot parameters The skip-lif-migration- before-reboot parameter specifies that LIF migration prior to a reboot will be skipped.		
Manually trigger a core dump when you shut down a node	system node halt with both the -dump and -skip-lif-migration-before-shutdown parametersThe skip-lif-migration- before-shutdown parameter specifies that LIF migration prior to a shutdown will be skipped.		
Save a specified core dump	system node coredump save		
Save all unsaved core dumps that are on a specified node	system node coredump save-all		

If you want to	Use this command
Generate and send an AutoSupport message with a core dump file you specify	system node autosupport invoke-core- upload The -uri optional parameter specifies an alternate destination for the
	AutoSupport message.
Display status information about core dumps	system node coredump status
Delete a specified core dump	system node coredump delete
Delete all unsaved core dumps or all saved core files on a node	system node coredump delete-all
Display application core dump reports	system node coredump reports show
Delete an application core dump report	system node coredump reports delete

Related information

ONTAP command reference

Disk and tier management

Disks and ONTAP local tiers

Local tiers, also called *aggregates*, are logical containers for the disks managed by a node. You can use local tiers to isolate workloads with different performance demands, to tier data with different access patterns, or to segregate data for regulatory purposes.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*.

- For business-critical applications that need the lowest possible latency and the highest possible performance, you might create a local tier consisting entirely of SSDs.
- To tier data with different access patterns, you can create a *hybrid local tier*, deploying flash as highperformance cache for a working data set, while using lower-cost HDDs or object storage for less frequently accessed data.
 - A Flash Pool consists of both SSDs and HDDs.
 - A *FabricPool* consists of an all-SSD local tier with an attached object store.
- If you need to segregate archived data from active data for regulatory purposes, you can use a local tier consisting of capacity HDDs, or a combination of performance and capacity HDDs.



You can use a FabricPool to tier data with different access patterns, deploying SSDs for frequently accessed "hot" data and object storage for rarely accessed "cold" data.

Working with local tiers in a MetroCluster configuration

If you have a MetroCluster configuration, you should following the procedures in the MetroCluster documentation for initial configuration and guidelines for local tiers and disk management.

Related information

- Manage local tiers
- Manage disks
- Manage RAID configurations
- Manage Flash Pool tiers
- Manage FabricPool cloud tiers

Manage local tiers

Learn about ONTAP local tier management

You can use System Manager or the ONTAP CLI to add local tiers, manage their usage, and add capacity (disks) to them.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

You can perform the following tasks:

• Add (create) a local tier

To add a local tier, you follow a specific workflow. You determine the number of disks or disk partitions that you need for the local tier and decide which method to use to create the local tier. You can add local tiers

automatically by letting ONTAP assign the configuration, or you can manually specify the configuration.

• Manage the use of local tiers

For existing local tiers, you can rename them, set their media costs, or determine their drive and RAID group information. You can modify the RAID configuration of a local tier and assign local tiers to storage VMs (SVMs).

You can modify the RAID configuration of a local tier and assign local tiers to storage VMs (SVMs). You can determine which volumes reside on a local tier and how much space they use on a local tier. You can control how much space that volumes can use. You can relocate local tier ownership with an HA pair. You can also delete a local tier.

• Add capacity (disks) to a local tier

Using different methods, you follow a specific workflow to add capacity. You can add disks to a local tier and add drives to a node or shelf. If needed, you can correct misaligned spare partitions.

Add (create) a local tier

Workflow to add an ONTAP local tier

Creating local tiers provides storage to volumes on your system.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

The workflow for creating local tiers is specific to the interface you use: System Manager or the CLI.

System Manager

System Manager creates local tiers based on recommended best practices for configuring local tiers.

Beginning with ONTAP 9.11.1, you can decide to configure local tiers manually if you want a different configuration than the one recommended during the automatic process to add a local tier.





Determine the number of disks or disk partitions required for an ONTAP local tier

You must have enough disks or disk partitions in your local tier to meet system and business requirements. You should also have the recommended number of hot spare disks or hot spare disk partitions to minimize the potential of data loss.

Root-data partitioning is enabled by default on certain configurations. Systems with root-data partitioning enabled use disk partitions to create local tiers. Systems that do not have root-data partitioning enabled use unpartitioned disks.

You must have enough disks or disk partitions to meet the minimum number required for your RAID policy and enough to meet your minimum capacity requirements.



In ONTAP, the usable space of the drive is less than the physical capacity of the drive. You can find the usable space of a specific drive and the minimum number of disks or disk partitions required for each RAID policy in the Hardware Universe.

Determine usable space of a specific disk

The procedure you follow depends on the interface you use-System Manager or the CLI:

System Manager

Use System Manager to determine usable space of disks

Perform the following steps to view the usable size of a disk:

Steps

- 1. Go to Storage > Tiers
- 2. Click inext to the name of the local tier.
- 3. Select the Disk Information tab.

CLI

Use the CLI to determine usable space of disks

Perform the following step to view the usable size of a disk:

Step

1. Display spare disk information:

```
storage aggregate show-spare-disks
```

In addition to the number of disks or disk partitions necessary to create your RAID group and meet your capacity requirements, you should also have the minimum number of hot spare disks or hot spare disk partitions recommended for your local tier:

• For all flash local tiers, you should have a minimum of one hot spare disk or disk partition.



The AFF C190 defaults to no spare drive. This exception is fully supported.

- For non-flash homogenous local tiers, you should have a minimum of two hot spare disks or disk partitions.
- For SSD storage pools, you should have a minimum of one hot spare disk for each HA pair.
- For Flash Pool local tiers, you should have a minimum of two spare disks for each HA pair. You can find more information on the supported RAID policies for Flash Pool local tiers in the Hardware Universe.
- To support the use of the Maintenance Center and to avoid issues caused by multiple concurrent disk failures, you should have a minimum of four hot spares in multi-disk carriers.

Related information

- NetApp Hardware Universe
- NetApp Technical Report 3838: Storage Subsystem Configuration Guide
- storage aggregate show

Decide which method to use to create ONTAP local tiers

Although ONTAP provides best-practice recommendations for adding local tiers automatically, you must determine whether the recommended configurations are supported in your environment. If they are not, you must make decisions about RAID policy and disk configuration and then create the local tiers manually. When a local tier is created automatically, ONTAP analyzes available spare disks in the cluster and generates a recommendation about how spare disks should be used to add local tiers according to best practices. ONTAP displays the recommended configurations. You can accept the recommendations or add the local tiers manually.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

Before you can accept ONTAP recommendations

If any of the following disk conditions are present, they must be addressed before accepting the recommendations from ONTAP:

- Missing disks
- · Fluctuation in spare disk numbers
- Unassigned disks
- · Non-zeroed spares
- Disks undergoing maintenance testing

When you must use the manual method

In many cases, the recommended layout of the local tier will be optimal for your environment. However, if your environment includes the following configurations, you must create the local tier using the manual method.



Beginning with ONTAP 9.11.1, you can manually add local tiers with System Manager.

- · Local tiers using third-party array LUNs
- Virtual disks with Cloud Volumes ONTAP or ONTAP Select
- MetroCluster system
- SyncMirror
- MSATA disks
- · Flash Pool tiers
- · Multiple disk types or sizes are connected to the node

Select the method to create local tiers

Choose which method you want to use:

- · Add (create) local tiers automatically
- Add (create) local tiers manually

Related information

- ONTAP command reference
- storage aggregate auto-provision

If the best-practice recommendation that ONTAP provides for automatically adding a local tier is appropriate in your environment, you can accept the recommendation and let ONTAP add the local tier.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

Before you begin

Disks must be owned by a node before they can be used in a local tier. If your cluster is not configured to use automatic disk ownership assignment, you must assign ownership manually.

Learn more about the commands described in this procedure in the ONTAP command reference.

System Manager

Steps

- 1. In System Manager, click **Storage > Tiers**.
- 2. From the **Tiers** page, click + Add Local Tier to create a new local tier:

The **Add Local Tier** page shows the recommended number of local tiers that can be created on the nodes and the usable storage available.

3. Click Recommended details to view the configuration recommended by System Manager.

System Manager displays the following information beginning with ONTAP 9.8:

- Local tier name (you can edit the local tier name beginning with ONTAP 9.10.1)
- Node name
- Usable size
- Type of storage

Beginning with ONTAP 9.10.1, additional information is displayed:

- · Disks: showing the number, size, and type of the disks
- **Layout**: showing the RAID group layout, including which disks are parity or data and which slots are unused.
- Spare disks: showing the node name, the number and size of spare disks, and the type of storage.
- 4. Perform one of the following steps:

If you want to	Then do this
Accept the recommendations from System Manager.	Proceed to the step for configuring the Onboard Key Manager for encryption.
Manually configure the local tiers and <i>not</i> use the recommendations from System Manager.	 Proceed to Add a local tier manually: For ONTAP 9.10.1 and earlier, follow the steps to use the CLI. Beginning with ONTAP 9.11.1, follow the steps to use System Manager.

- 5. (Optional): If the Onboard Key Manager has been installed, you can configure it for encryption. Check the **Configure Onboard Key Manager for encryption** check box.
 - a. Enter a passphrase.
 - b. Enter the passphrase again to confirm it.
 - c. Save the passphrase for future use in case the system needs to be recovered.
 - d. Back up the key database for future use.
- 6. Click Save to create the local tier and add it to your storage solution.

You run the storage aggregate auto-provision command to generate local tier layout recommendations. You can then create local tiers after reviewing and approving ONTAP recommendations.

About this task

The default summary generated with the storage aggregate auto-provision command lists the recommended local tiers to be created, including names and usable size. You can view the list and determine whether you want to create the recommended local tiers when prompted.

You can also display a detailed summary by using the -verbose option, which displays the following reports:

- Per node summary of new local tiers to create, discovered spares, and remaining spare disks and partitions after local tier creation
- New data local tiers to create with counts of disks and partitions to be used
- RAID group layout showing how spare disks and partitions will be used in new data local tiers to be created
- Details about spare disks and partitions remaining after local tier creation

If you are familiar with the auto-provision method and your environment is correctly prepared, you can use the -skip-confirmation option to create the recommended local tier without display and confirmation. The storage aggregate auto-provision command is not affected by the CLI session -confirmations setting.

Learn more about storage aggregate auto-provision in the ONTAP command reference.

Steps

- 1. Run the storage aggregate auto-provision command with the desired display options.
 - no options: Display standard summary
 - -verbose option: Display detailed summary
 - -skip-confirmation option: Create recommended local tiers without display or confirmation
- 2. Perform one of the following steps:

If you want to	Then do this
----------------	--------------

Accept the recommendations from ONTAP.	Review the display of recommended local tiers, and then respond to the prompt to create the recommended local tiers.			
	myA400-44556677::> storage aggregate auto- provision Node New Data Aggregate Usable Size			
	myA400-364 myA400_364_SSD_1 3.29TB			
	myA400-363 myA400_363_SSD_1 1.46TB 			
	Total: 2 new data aggregates 4.75TB			
	Do you want to create recommended aggregates? {y n}: y			
	Info: Aggregate auto provision has started. Use the "storage aggregate show-auto-provision-progress" command to track the progress.			
	myA400-44556677::>			
Manually configure the local tiers and <i>not</i> use the recommendations from ONTAP.	Proceed to Add a local tier manually.			

Related information

• ONTAP command reference

Add ONTAP local tiers manually

If you do not want to add a local tier using the best-practice recommendations from ONTAP, you can perform the process manually.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see <u>Disks and local tiers</u>.

Before you begin

Disks must be owned by a node before they can be used in a local tier. If your cluster is not configured to use automatic disk ownership assignment, you must assign ownership manually.

Learn more about the commands described in this procedure in the ONTAP command reference.

System Manager

Beginning with ONTAP 9.11.1, if you do not want to use the configuration recommended by System Manager to create a local tier, you can specify the configuration you want.

Steps

- 1. In System Manager, click **Storage > Tiers**.
- 2. From the **Tiers** page, click + Add Local Tier to create a new local tier:

The **Add Local Tier** page shows the recommended number of local tiers that can be created on the nodes and the usable storage available.

3. When System Manager displays the storage recommendation for the local tier, click **Switch to Manual Local Tier Creation** in the **Spare Disks** section.

The Add Local Tier page displays fields that you use to configure the local tier.

- 4. In the first section of the Add Local Tier page, complete the following:
 - a. Enter the name of the local tier.
 - b. (Optional): Check the Mirror this local tier check box if you want to mirror the local tier.
 - c. Select a disk type.
 - d. Select the number of disks.
- 5. In the RAID Configuration section, complete the following:
 - a. Select the RAID type.
 - b. Select the RAID group size.
 - c. Click RAID allocation to view how the disks are allocated in the group.
- 6. (Optional): If the Onboard Key Manager has been installed, you can configure it for encryption in the **Encryption** section of the page. Check the **Configure Onboard Key Manager for encryption** check box.
 - a. Enter a passphrase.
 - b. Enter the passphrase again to confirm it.
 - c. Save the passphrase for future use in case the system needs to be recovered.
 - d. Back up the key database for future use.
- 7. Click **Save** to create the local tier and add it to your storage solution.

CLI

Before you create local tiers manually, you should review disk configuration options and simulate creation.

Then you can issue the storage aggregate create command and verify the results.

Before you begin

You must have determined the number of disks and the number of hot spare disks you need in the local tier.

About this task

If root-data-data partitioning is enabled and you have 24 solid-state drives (SSDs) or fewer in your configuration, it is recommended that your data partitions be assigned to different nodes.

The procedure for creating local tiers on systems with root-data partitioning and root-data-data partitioning enabled is the same as the procedure for creating local tiers on systems using unpartitioned disks. If root-data partitioning is enabled on your system, you should use the number of disk partitions for the -diskcount option. For root-data-data partitioning, the -diskcount option specifies the count of disks to use.



When creating multiple local tiers for use with FlexGroup volumes, local tiers should be as close in size as possible.

Learn more about storage aggregate create and local tier creation options and requirements in the ONTAP command reference.

Steps

1. View the list of spare disk partitions to verify that you have enough to create your local tier:

storage aggregate show-spare-disks -original-owner node name

Data partitions are displayed under Local Data Usable. A root partition cannot be used as a spare.

2. Simulate the creation of the local tier:

```
storage aggregate create -aggregate aggregate_name -node node_name
-raidtype raid dp -diskcount number of disks or partitions -simulate true
```

- 3. If any warnings are displayed from the simulated command, adjust the command and repeat the simulation.
- 4. Create the local tier:

```
storage aggregate create -aggregate aggr_name -node node_name -raidtype
raid_dp -diskcount number_of_disks_or_partitions
```

5. Display the local tier to verify that it was created:

storage aggregate show-status aggregate name

Related information

• storage aggregate show

Add ONTAP local tiers with SyncMirror enabled

You can enable SyncMirror when you manually create a local tier to synchronously mirror local tier data.

Learn more about mirrored and unmirrored local tiers.

Before you begin

- The cluster must be initialized with only internal storage.
- Cluster setup must have been completed on both nodes.

About this task

This procedure creates mirrored data local tiers of equal size on each cluster node, and each local tier has a disk count of 44.

Steps

1. Disable storage auto-assignment:

storage disk option modify -node * -autoassign off

2. Confirm that auto-assignment is disabled:

storage disk option show

- 3. Attach the external shelf.
- 4. Assign the external drives to each node specifying pool 1:

storage disk assign -disk <disk ID> -owner <node name> -pool 1

5. Mirror the root local tier on each node:

storage aggregate mirror -aggregate <node1 root-aggr>

storage aggregate mirror -aggregate <node2 root-aggr>



Drives in pool 1 are automatically partitioned to match those in plex 0.

6. On node 1, create a mirrored data local tier using a disk count of 44. This selects 22 partitions from pool 0 and 22 partitions from pool 1.

storage aggregate create -node <nodel name> -aggregate <nodel aggr-name>
-diskcount 44 -mirror true

7. On node 2, create a mirrored data local tier using a disk count of 44. This selects 22 partitions from pool 0 and 22 partitions from pool 1.

storage aggregate create -node <node2 name> -aggregate <node2 aggr-name>
-diskcount 44 -mirror true

8. Verify that local tiers of equal size were successfully created:

storage aggregate show

Manage the use of local tiers

Rename an ONTAP local tier

You can rename a local tier. The method you follow depends on the interface you use— System Manager or the CLI.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

System Manager

Use System Manager to rename a local tier

Beginning with ONTAP 9.10.1, you can modify the name of a local tier.

Steps

- 1. In System Manager, click **Storage > Tiers**.
- 2. Click next to the name of the local tier.
- 3. Select Rename.
- 4. Specify a new name for the local tier.

CLI

Use the CLI to rename a local tier

Step

1. Using the CLI, rename the local tier:

storage aggregate rename -aggregate aggr-name -newname aggr-new-name

The following example renames an aggregate named "aggr5" as "sales-aggr":

> storage aggregate rename -aggregate aggr5 -newname sales-aggr

Related information

• storage aggregate rename

Set media cost of an ONTAP local tier

Beginning with ONTAP 9.11.1, you can use System Manager to set the media cost of a local tier.

Steps

- 1. In System Manager, click **Storage > Tiers**, then click **Set Media Cost** in the desired local tier tiles.
- 2. Select active and inactive tiers to enable comparison.
- 3. Enter a currency type and amount.

When you enter or change the media cost, the change is made in all media types.

Manually fast zero ONTAP drives

On systems freshly installed with ONTAP 9.4 or later and systems reinitialized with ONTAP 9.4 or later, *fast zeroing* is used to zero drives.

With *fast zeroing*, drives are zeroed in seconds. This is done automatically before provisioning and greatly reduces the time it takes to initialize the system, create local tiers, or expand local tiers when spare drives are added.

Fast zeroing is supported on both SSDs and HDDs.



Fast zeroing is not supported on systems upgraded from ONTAP 9.3 or earlier. ONTAP 9.4 or later must be freshly installed or the system must be reinitialized. In ONTAP 9.3 and earlier, drives are also automatically zeroed by ONTAP; however, the process takes longer.

If you need to manually zero a drive, you can use one of the following methods. In ONTAP 9.4 and later, manually zeroing a drive also takes only seconds.

CLI command

Use a CLI command to fast-zero drives

About this task

Admin privileges are required to use this command.

Steps

1. Enter the CLI command:

storage disk zerospares

Boot menu options

Select options from the boot menu to fast-zero drives

About this task

• The fast zeroing enhancement does not support systems upgraded from a release earlier than ONTAP 9.4.

Steps

- 1. From the boot menu, select one of the following options:
 - (4) Clean configuration and initialize all disks
 - (9a) Unpartition all disks and remove their ownership information
 - \circ (9b) Clean configuration and initialize node with whole disks

Manually assign ONTAP disk ownership

Disks must be owned by a node before they can be used in a local tier.

About this task

- If you are manually assigning ownership in an HA pair that is not being initialized and does not have only DS460C shelves, use option 1.
- If you are initializing an HA pair that has only DS460C shelves, use option 2 to manually assign ownership for the root drives.

Option 1: Most HA pairs

For an HA pair that is not being initialized and does not have only DS460C shelves, use this procedure to manually assigning ownership.

About this task

- The disks you are assigning ownership for must be in a shelf that is physically cabled to the node you are assigning ownership to.
- If you are using disks in a local tier (aggregate):
 - Disks must be owned by a node before they can be used in a local tier (aggregate).
 - You cannot reassign ownership of a disk that is in use in a local tier (aggregate).

Steps

1. Use the CLI to display all unowned disks:

storage disk show -container-type unassigned

2. Assign each disk:

storage disk assign -disk disk name -owner owner name

You can use the wildcard character to assign more than one disk at once. If you are reassigning a spare disk that is already owned by a different node, you must use the "-force" option.

For an HA pair that you are initializing and that only has DS460C shelves, use this procedure to manually assign ownership for the root drives.

About this task

• When you initialize an HA pair that has only DS460C shelves, you must manually assign the root drives to conform to the half-drawer policy.

After HA pair initialization (boot up), automatic assignment of disk ownership is automatically enabled and uses the half-drawer policy to assign ownership to the remaining drives (other than the root drives) and any drives added in the future, such as replacing failed disks, responding to a "low spares" message, or adding capacity.

Learn about the half-drawer policy.

• RAID needs a minimum of 10 drives for each HA pair (5 for each node) for any greater than 8TB NL-SAS drives in a DS460C shelf.

Steps

- 1. If your DS460C shelves are not fully populated, complete the following substeps; otherwise, go to the next step.
 - a. First, install drives in the front row (drive bays 0, 3, 6, and 9) of each drawer.

Installing drives in the front row of each drawer allows for proper air flow and prevents overheating.

b. For the remaining drives, evenly distribute them across each drawer.

Fill drawer rows from front to back. If you don't have enough drives to fill rows, then install them in pairs so that drives occupy the left and right side of a drawer evenly.

The following illustration shows the drive bay numbering and locations in a DS460C drawer.



2. Log into the clustershell using the node-management LIF or cluster-management LIF.

3. Manually assign the root drives in each drawer to conform to the half-drawer policy using the following substeps:

The half-drawer policy has you assign the left half of a drawer's drives (bays 0 to 5) to node A, and the right half of a drawer's drives (bays 6 to 11) to node B.

- a. Display all unowned disks: storage disk show -container-type unassigned
- b. Assign the root disks: storage disk assign -disk disk name -owner owner name

You can use the wildcard character to assign more than one disk at a time.

Learn more about storage disk in the ONTAP command reference.

Determine drive and RAID group information for an ONTAP local tier

Some local tier administration tasks require that you know what types of drives compose the local tier, their size, checksum, and status, whether they are shared with other local tiers, and the size and composition of the RAID groups.

Step

1. Show the drives for the local tier, by RAID group:

storage aggregate show-status aggr_name

The drives are displayed for each RAID group in the local tier.

You can see the RAID type of the drive (data, parity, dparity) in the Position column. If the Position column displays shared, then the drive is shared: if it is an HDD, it is a partitioned disk; if it is an SSD, it is part of a storage pool.

```
cluster1::> storage aggregate show-status nodeA fp 1
Owner Node: cluster1-a
Aggregate: nodeA fp 1 (online, mixed raid type, hybrid) (block checksums)
 Plex: /nodeA_fp_1/plex0 (online, normal, active, pool0)
  RAID Group /nodeA fp 1/plex0/rg0 (normal, block checksums, raid dp)
                                     Usable Physical
    Position Disk Pool Type
                                      Size
                                              Size Status
                                RPM
    _____
                    _____ _____ _____ _____ _____ _____
    shared 2.0.1
                        SAS
                              10000 472.9GB 547.1GB (normal)
                     0
    shared 2.0.3
                    0 SAS
                             10000 472.9GB 547.1GB (normal)
                     0 SAS
                             10000 472.9GB 547.1GB (normal)
    shared 2.0.5
    shared 2.0.7
                     0 SAS
                             10000 472.9GB 547.1GB (normal)
    shared 2.0.9
                     0 SAS
                             10000 472.9GB 547.1GB (normal)
    shared 2.0.11
                     0 SAS 10000 472.9GB 547.1GB (normal)
  RAID Group /nodeA flashpool 1/plex0/rg1
  (normal, block checksums, raid4) (Storage Pool: SmallSP)
                                     Usable Physical
    Position Disk Pool Type RPM
                                      Size
                                              Size Status
    shared 2.0.13 0 SSD
                              - 186.2GB 745.2GB (normal)
    shared 2.0.12 0 SSD - 186.2GB 745.2GB (normal)
8 entries were displayed.
```

Related information

storage aggregate show-status

Assign ONTAP local tiers to storage VMs (SVMs)

If you assign one or more local tiers to a storage virtual machine (storage VM or SVM, formerly known as Vserver), then you can use only those local tiers to contain volumes for that storage VM (SVM).



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

Before you begin

The storage VM and the local tiers you want to assign to that storage VM must already exist.

About this task

Assigning local tiers to your storage VMs helps you keep your storage VMs isolated from each other; this is especially important in a multi-tenancy environment.

Steps

1. Check the list of local tiers already assigned to the SVM:

```
vserver show -fields aggr-list
```

The local tiers currently assigned to the SVM are displayed. If there are no local tiers assigned, – is displayed.

2. Add or remove assigned local tiers, depending on your requirements:

If you want to	Use this command
Assign additional local tiers	vserver add-aggregates
Unassign local tiers	vserver remove-aggregates

The listed local tiers are assigned to or removed from the SVM. If the SVM already has volumes that use an aggregate that is not assigned to the SVM, a warning message is displayed, but the command is completed successfully. Any local tiers that were already assigned to the SVM and that were not named in the command are unaffected.

Example

In the following example, the local tiers aggr1 and aggr2 are assigned to SVM svm1:

vserver add-aggregates -vserver svm1 -aggregates aggr1,aggr2

Determine which volumes reside on an ONTAP local tier

You might need to determine which volumes reside on a local tier before performing operations on the local tier, such as relocating it or taking it offline.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

Steps

1. To display the volumes that reside on a local tier, enter

volume show -aggregate aggregate_name

All volumes that reside on the specified local tier are displayed.

Determine and control space usage of a volume in an ONTAP local tier

You can determine which FlexVol volumes are using the most space in a local tier and specifically which features within the volume.

The volume show-footprint command provides information about a volume's footprint, or its space usage within the containing local tier.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

The volume show-footprint command shows details about the space usage of each volume in a local tier, including offline volumes. This command bridges the gap between the output of the volume show-space and aggregate show-space commands. All percentages are calculated as a percent of local tier size.

The following example shows the volume show-footprint command output for a volume called testvol:

```
cluster1::> volume show-footprint testvol
     Vserver : thevs
     Volume : testvol
     Feature
                                          Used
                                                  Used%
     _____
                                     _____
                                                  ____
     Volume Data Footprint
                                        120.6MB
                                                    48
     Volume Guarantee
                                        1.88GB
                                                   71%
     Flexible Volume Metadata
                                        11.38MB
                                                    0%
     Delayed Frees
                                         1.36MB
                                                    0%
     Total Footprint
                                         2.01GB
                                                    76%
```

The following table explains some of the key rows of the output of the volume show-footprint command and what you can do to try to decrease space usage by that feature:

Row/feature name	Description/contents of row	Some ways to decrease
Volume Data Footprint	The total amount of space used in the containing local tier by a volume's data in the active file system and the space used by the volume's snapshots. This row does not include reserved space.	 Deleting data from the volume. Deleting snapshots from the volume.
Volume Guarantee	The amount of space reserved by the volume in the local tier for future writes. The amount of space reserved depends on the guarantee type of the volume.	Changing the type of guarantee for the volume to none.
Flexible Volume Metadata	The total amount of space used in the local tier by the volume's metadata files.	No direct method to control.

Delayed Frees	Blocks that ONTAP used for performance and cannot be immediately freed. For SnapMirror destinations, this row has a value of 0 and is not displayed.	No direct method to control.
File Operation Metadata	The total amount of space reserved for file operation metadata.	No direct method to control.
Total Footprint	The total amount of space that the volume uses in the local tier. It is the sum of all of the rows.	Any of the methods used to decrease space used by a volume.

Related information

NetApp Technical Report 3483: Thin Provisioning in a NetApp SAN or IP SAN Enterprise Environment

Determine space usage in an ONTAP local tier

You can view how much space is used by all volumes in one or more local tiers so that you can take actions to free more space.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

WAFL reserves a percentage of the total disk space for local tier level metadata and performance. The space used for maintaining the volumes in the local tier comes out of the WAFL reserve and cannot be changed.

In local tiers smaller than 30 TB, WAFL reserves 10% of the total disk space for local tier level metadata and performance.

Beginning with ONTAP 9.12.1, in local tiers that are 30 TB or larger, the amount of reserved disk space for local tier level metadata and performance is reduced, resulting in 5% more usable space in local tiers. The availability of this space savings varies based on your platform and version of ONTAP.

Disk space reserved by ONTAP in local tiers 30 TB or greater	Applies to platforms	In ONTAP versions
5%	All AFF and FAS platforms	ONTAP 9.14.1 and later
5%	AFF platforms and FAS500f platforms	ONTAP 9.12.1 and later
10%	All platforms	ONTAP 9.11.1 and later

You can view space usage by all volumes in one or more local tiers with the aggregate show-space command. This helps you see which volumes are consuming the most space in their containing local tiers so that you can take actions to free more space.

The used space in an local tier is directly affected by the space used in the FlexVol volumes it contains. Measures that you take to increase space in a volume also affect space in the local tier.



Beginning with ONTAP 9.15.1, two new metadata counters are available. Together with changes to several existing counters, you can get a clearer view of the amount of user data allocated. See Determine space usage in a volume or local tier for more information.

The following rows are included in the aggregate show-space command output:

Volume Footprints

The total of all volume footprints within the local tier. It includes all of the space that is used or reserved by all data and metadata of all volumes in the containing local tier.

Aggregate Metadata

The total file system metadata required by the local tier, such as allocation bitmaps and inode files.

Snapshot Reserve

The amount of space reserved for local tier snapshots, based on volume size. It is considered used space and is not available to volume or local tier data or metadata.

Snapshot Reserve Unusable

The amount of space originally allocated for local tier snapshot reserve that is unavailable for local tier snapshots because it is being used by volumes associated with the local tier. Can occur only for local tiers with a non-zero local tier snapshot reserve.

Total Used

The sum of all space used or reserved in the local tier by volumes, metadata, or snapshots.

Total Physical Used

The amount of space being used for data now (rather than being reserved for future use). Includes space used by local tier snapshots.

The following example shows the aggregate show-space command output for an local tier whose snapshot reserve is 5%. If the snapshot reserve was 0, the row would not be displayed.

<pre>cluster1::> storage aggregate show-space</pre>				
Aggregate : wqa_gx106_aggr1				
Feature	Used	Used%		
Volume Footprints Aggregate Metadata Snapshot Reserve	101.0MB 300KB 5.98GB	 0% 5%		
Total Used Total Physical Used	6.07GB 34.82KB	5% 0%		

Related Information

- Knowledge Base article: Space Usage
- Free up 5% of your storage capacity by upgrading to ONTAP 9.12.1
- storage aggregate show-space

Relocate ownership of an ONTAP local tier within an HA pair

You can change the ownership of local tiers among the nodes in an HA pair without interrupting service from the local tiers.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

Both nodes in an HA pair are physically connected to each other's disks or array LUNs. Each disk or array LUN is owned by one of the nodes.

Ownership of all disks or array LUNs within a local tier changes temporarily from one node to the other when a takeover occurs. However, local tiers relocation operations can also permanently change the ownership (for example, if done for load balancing). The ownership changes without any data-copy processes or physical movement of the disks or array LUNs.

About this task

• Because volume count limits are validated programmatically during local tier relocation operations, it is not necessary to check for this manually.

If the volume count exceeds the supported limit, the local tier relocation operation fails with a relevant error message.

• You should not initiate local tier relocation when system-level operations are in progress on either the source or the destination node; likewise, you should not start these operations during the local tier relocation.

These operations can include the following:

- Takeover
- Giveback
- Shutdown
- · Another local tier relocation operation
- Disk ownership changes
- Local tier or volume configuration operations
- · Storage controller replacement
- ONTAP upgrade
- ONTAP revert
- If you have a MetroCluster configuration, you should not initiate local tier relocation while disaster recovery operations (*switchover, healing*, or *switchback*) are in progress.
- If you have a MetroCluster configuration and initiate local tier relocation on a switched-over local tier, the operation might fail because it exceeds the DR partner's volume limit count.

- You should not initiate local tier relocation on local tiers that are corrupt or undergoing maintenance.
- Before initiating the local tier relocation, you should save any core dumps on the source and destination nodes.

Steps

1. View the local tiers on the node to confirm which local tiers to move and ensure they are online and in good condition:

storage aggregate show -node *source-node*

The following command shows six local tiers on the four nodes in the cluster. All local tiers are online. Node1 and Node3 form an HA pair and Node2 and Node4 form an HA pair.

cluster::> Aggregate	storage Size	aggregate Available	show Used%	State	#Vols	Nodes	RAID Status
aqqr 0	239.0GB	11.13GB	 95%	online		nodel	raid dp,
							normal
aggr_1	239.0GB	11.13GB	95%	online	1	node1	raid_dp,
aggr_2	239.0GB	11.13GB	95%	online	1	node2	raid_dp,
aggr_3	239.0GB	11.13GB	95%	online	1	node2	normal raid_dp,
			•		_		normal
aggr_4	239.0GB	238.9GB	08	online	5	node3	raid_dp,
aggr_5	239.0GB	239.0GB	0%	online	4	node4	raid_dp,
6 entries	were disp	blayed.					normar

2. Issue the command to start the local tier relocation:

storage aggregate relocation start -aggregate-list aggregate-1, aggregate-2...
-node source-node -destination destination-node

The following command moves the local tier aggr_1 and aggr_2 from Node1 to Node3. Node3 is Node1's HA partner. The local tiers can be moved only within the HA pair.

cluster::> storage aggregate relocation start -aggregate-list aggr_1, aggr_2 -node node1 -destination node3 Run the storage aggregate relocation show command to check relocation status. node1::storage aggregate>

3. Monitor the progress of the local tier relocation with the storage aggregate relocation show command:

```
storage aggregate relocation show -node source-node
```

The following command shows the progress of the local tiers that are being moved to Node3:

When the relocation is complete, the output of this command shows each local tier with a relocation status of "Done".

Related information

- storage aggregate relocation show
- storage aggregate relocation start
- storage aggregate show

Delete an ONTAP local tier

You can delete a local tier if there are no volumes on the local tier.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

The storage aggregate delete command deletes a storage local tier. The command fails if there are volumes present on the local tier. If the local tier has an object store attached to it, then in addition to deleting the local, the command deletes the objects in the object store as well. No changes are made to the object store configuration as part of this command.

The following example deletes an local tier named "aggr1":

> storage aggregate delete -aggregate aggr1

Related information

storage aggregate delete

ONTAP commands for a local tier relocation

There are specific ONTAP commands for relocating local tier ownership within an HA pair.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

If you want to	Use this command
Start the local tier relocation process	storage aggregate relocation start
Monitor the local tier relocation process	storage aggregate relocation show

Related information

- storage aggregate relocation show
- storage aggregate relocation start

ONTAP commands for managing local tiers

You use the storage aggregate command to manage your local tiers.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

If you want to	Use this command
Display the size of the cache for all Flash Pool local tiers	storage aggregate show -fields hybrid- cache-size-total -hybrid-cache-size -total >0
Display disk information and status for an local tier	storage aggregate show-status
Display spare disks by node	storage aggregate show-spare-disks
Display the root local tiers in the cluster	storage aggregate show -has-mroot true
Display basic information and status for local tiers	storage aggregate show
Display the type of storage used in an local tiers	storage aggregate show -fields storage- type
Bring an local tier online	storage aggregate online
Delete an local tier	storage aggregate delete
Put an local tier into the restricted state	storage aggregate restrict

If you want to	Use this command
Rename an local tier	storage aggregate rename
Take an local tier offline	storage aggregate offline
Change the RAID type for a local tier	storage aggregate modify -raidtype

Related information

- storage aggregate delete
- storage aggregate modify
- storage aggregate offline
- storage aggregate online
- storage aggregate rename
- storage aggregate restrict
- storage aggregate show

Add capacity (disks) to a local tier

Workflow to add capacity to an ONTAP local tier

To add capacity to a local tier you must first identify which local tier you want to add to, determine how much new storage is needed, install new disks, assign disk ownership, and create a new RAID group, if needed.

You can use either System Manager or the ONTAP CLI to add capacity.



Methods to create space in an ONTAP local tier

If a local tier runs out of free space, various problems can result that range from loss of data to disabling a volume's guarantee. There are multiple ways to make more space in a local tier.

All of the methods have various consequences. Prior to taking any action, you should read the relevant section in the documentation.

The following are some common ways to make space in local tier, in order of least to most consequences:

- Add disks to the local tier.
- Move some volumes to another local tier with available space.
- Shrink the size of volume-guaranteed volumes in the local tier.
- Delete unneeded volume snapshots if the volume's guarantee type is "none".
- Delete unneeded volumes.
- Enable space-saving features, such as deduplication or compression.
- (Temporarily) disable features that are using a large amount of metadata .

Add capacity to an ONTAP local tier

You can add disks to an local tier so that it can provide more storage to its associated volumes.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

System Manager (ONTAP 9.8 and later)



Beginning with ONTAP 9.12.1, you can use System Manager to view the committed capacity of a local tier to determine if additional capacity is required for the local tier. See Monitor capacity in System Manager.

Steps

- 1. Select **Storage > Tiers**.
- 2. Select next to the name of the local tier to which you want to add capacity.
- 3. Select Add Capacity.



If there are no spare disks that you can add, then the **Add Capacity** option is not shown, and you cannot increase the capacity of the local tier.

4. Perform the following steps, based on the version of ONTAP that is installed:

If this version of ONTAP is installed	Perform these steps
Beginning with ONTAP 9.11.1	 Select the disk type and number of disks. If you want to add disks to a new RAID group, check the check box. The RAID allocation is displayed. Select Save.
ONTAP 9.10.1, 9.9, or 9.8	 If the node contains multiple storage tiers, then select the number of disks you want to add to the local tier. Otherwise, if the node contains only a single storage tier, the added capacity is estimated automatically. Select Add.

- (Optional) The process takes some time to complete. If you want to run the process in the background, select **Run in Background**.
- 6. After the process completes, you can view the increased capacity amount in the local tier information at **Storage > Tiers**.

System Manager (ONTAP 9.7 and earlier)

Steps

- 1. (For ONTAP 9.7 only) Select (Return to classic version).
- 2. Select Hardware and Diagnostics > Aggregates.
- 3. Select the local tier to which you want to add capacity disks, and then select **Actions > Add Capacity**.



You should add disks that are of the same size as the other disks in the local tier.

- 4. (For ONTAP 9.7 only) Select Switch to the new experience.
- 5. Select **Storage > Tiers** to verify the size of the new local tier.

CLI

Before you begin

You must know what the RAID group size is for the local tier you are adding the storage to.

About this task

This procedure for adding partitioned disks to a local tier is similar to the procedure for adding unpartitioned disks.

When you expand a local tier, you should be aware of whether you are adding partition or unpartitioned disks to the local tier. When you add unpartitioned drives to an existing local tier, the size of the existing RAID groups is inherited by the new RAID group, which can affect the number of parity disks required. If an unpartitioned disk is added to a RAID group composed of partitioned disks, the new disk is partitioned, leaving an unused spare partition.

When you provision partitions, you must ensure that you do not leave the node without a drive with both partitions as spare. If you do, and the node experiences a controller disruption, valuable information about the problem (the core file) might not be available to provide to the technical support.

Steps

1. Show the available spare storage on the system that owns the local tier:

```
storage aggregate show-spare-disks -original-owner node name
```

You can use the *-is-disk-shared* parameter to show only partitioned drives or only unpartitioned drives.

cl1-s2::> storage aggregate show-spare-disks -original-owner cl1-s2 -is-disk-shared true Original Owner: cl1-s2 PoolO Shared HDD Spares Local Local Data Root Physical Disk Type RPM Checksum Usable Usable Size Status _____ ____ ----- ----- ------1.0.1 BSAS 7200 block 753.8GB 73.89GB 828.0GB zeroed 1.0.2 BSAS 7200 block 753.8GB 0B 828.0GB zeroed 1.0.3 753.8GB BSAS 7200 block OB 828.0GB zeroed BSAS 7200 block 1.0.4 753.8GB 0B 828.0GB zeroed BSAS 7200 block 753.8GB 1.0.8 OB 828.0GB zeroed 1.0.9 BSAS 7200 block 753.8GB OB 828.0GB zeroed 1.0.10 BSAS 7200 block 0в 73.89GB 828.0GB zeroed 2 entries were displayed.

2. Show the current RAID groups for the local tier:

storage aggregate show-status <aggr name>
```
cl1-s2::> storage aggregate show-status -aggregate data 1
Owner Node: cl1-s2
Aggregate: data 1 (online, raid dp) (block checksums)
 Plex: /data 1/plex0 (online, normal, active, pool0)
  RAID Group /data 1/plex0/rg0 (normal, block checksums)
                                     Usable Physical
   Position Disk Pool Type RPM Size Size Status
    _____
                                  _ _____ ____
_____
    shared 1.0.10 0 BSAS 7200 753.8GB 828.0GB
(normal)
    shared 1.0.5 0 BSAS 7200 753.8GB 828.0GB
(normal)
    shared 1.0.6 0 BSAS 7200 753.8GB 828.0GB
(normal)
    shared 1.0.11 0 BSAS 7200 753.8GB 828.0GB
(normal)
    shared 1.0.0 0 BSAS 7200 753.8GB 828.0GB
(normal)
5 entries were displayed.
```

3. Simulate adding the storage to the aggregate:

```
storage aggregate add-disks -aggregate <aggr_name> -diskcount
<number of disks or partitions> -simulate true
```

You can see the result of the storage addition without actually provisioning any storage. If any warnings are displayed from the simulated command, you can adjust the command and repeat the simulation.

```
cl1-s2::> storage aggregate add-disks -aggregate aggr test
-diskcount 5 -simulate true
Disks would be added to aggregate "aggr test" on node "cl1-s2" in
the
following manner:
First Plex
 RAID Group rg0, 5 disks (block checksum, raid dp)
                                                 Usable
Physical
   Position Disk
                                     Туре
                                                   Size
Size
   _____
                                     _____
_____
   shared 1.11.4
                                     SSD
                                                415.8GB
415.8GB
   shared 1.11.18
                                     SSD
                                                415.8GB
415.8GB
   shared
            1.11.19
                                     SSD
                                                415.8GB
415.8GB
                                                415.8GB
   shared 1.11.20
                                     SSD
415.8GB
   shared 1.11.21
                                                415.8GB
                                     SSD
415.8GB
Aggregate capacity available for volume use would be increased by
1.83TB.
```

4. Add the storage to the aggregate:

```
storage aggregate add-disks -aggregate <aggr_name> -raidgroup new
-diskcount <number_of_disks_or_partitions>
```

When creating a Flash Pool local tier, if you are adding disks with a different checksum than the local tier, or if you are adding disks to a mixed checksum local tier, you must use the -checksumstyle parameter.

If you are adding disks to a Flash Pool local tier, you must use the -disktype parameter to specify the disk type.

You can use the -disksize parameter to specify a size of the disks to add. Only disks with approximately the specified size are selected for addition to the local tier.

```
cl1-s2::> storage aggregate add-disks -aggregate data_1 -raidgroup
new -diskcount 5
```

5. Verify that the storage was added successfully:

```
storage aggregate show-status -aggregate <aggr name>
```

```
cl1-s2::> storage aggregate show-status -aggregate data 1
Owner Node: cl1-s2
Aggregate: data 1 (online, raid dp) (block checksums)
 Plex: /data 1/plex0 (online, normal, active, pool0)
  RAID Group /data_1/plex0/rg0 (normal, block checksums)
                                                        Usable
Physical
    Position Disk
                                      Pool Type RPM
                                                         Size
Size Status
    _____
_____ ___
    shared 1.0.10
                                       0
                                           BSAS 7200 753.8GB
828.0GB (normal)
    shared 1.0.5
                                       0
                                           BSAS 7200 753.8GB
828.0GB (normal)
    shared 1.0.6
                                          BSAS 7200 753.8GB
                                       0
828.0GB (normal)
    shared 1.0.11
                                       0
                                           BSAS 7200 753.8GB
828.0GB (normal)
    shared 1.0.0
                                       0
                                           BSAS 7200 753.8GB
828.0GB (normal)
    shared 1.0.2
                                       0
                                           BSAS 7200 753.8GB
828.0GB (normal)
                                          BSAS 7200 753.8GB
    shared 1.0.3
                                       0
828.0GB (normal)
    shared 1.0.4
                                       0
                                          BSAS 7200 753.8GB
828.0GB (normal)
                                           BSAS 7200 753.8GB
    shared 1.0.8
                                       0
828.0GB (normal)
    shared 1.0.9
                                       0
                                         BSAS 7200 753.8GB
828.0GB (normal)
10 entries were displayed.
```

6. Verify that the node still has at least one drive with both the root partition and the data partition as spare:

storage aggregate show-spare-disks -original-owner <node name>

```
cl1-s2::> storage aggregate show-spare-disks -original-owner cl1-s2
-is-disk-shared true
Original Owner: cl1-s2
 PoolO
 Shared HDD Spares
                                                  Local
Local
                                                   Data
Root Physical
Disk
                        Type RPM Checksum
                                                Usable
Usable Size Status
 _____
                                      _____ ____
_____ _ ___
                        BSAS 7200 block
1.0.1
                                              753.8GB
73.89GB 828.0GB zeroed
1.0.10
                        BSAS 7200 block
                                                     0B
73.89GB 828.0GB zeroed
2 entries were displayed.
```

Related information

- storage aggregate add-disks
- storage aggregate show-spare-disks
- storage aggregate show-status

Add drives to an ONTAP node or shelf

You add drives to a node or shelf to increase the number of hot spares or to add space to local tier.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

Before you begin

The drive you want to add must be supported by your platform. You can confirm using the NetApp Hardware Universe.

The minimum number of drives you should add in a single procedure is six. Adding a single drive might reduce performance.

Steps for the NetApp Hardware Universe

- 1. In the Products dropdown menu, select your hardware configuration
- 2. Select your platform.
- 3. Select the version of ONTAP you are running then **Show Results**.
- 4. Beneath the graphic, select **Click here to see alternate views**. Choose the view that matches your configuration.



Steps to install the drives

1. Check the NetApp Support Site for newer drive and shelf firmware and Disk Qualification Package files.

If your node or shelf does not have the latest versions, update them before installing the new drive.

Drive firmware is automatically updated (nondisruptively) on new drives that do not have current firmware versions.

- 2. Properly ground yourself.
- 3. Gently remove the bezel from the front of the platform.
- 4. Identify the correct slot for the new drive.



The correct slots for adding drives vary depending on the platform model and ONTAP version. In some cases you need to add drives to specific slots in sequence. For example, in an AFF A800 you add the drives at specific intervals leaving clusters of empty slots. Whereas, in an AFF A220 you add new drives to the next empty slots running from the outside towards the middle of the shelf.

Refer to the steps in **Before you begin** to identify the correct slots for your configuration in the NetApp Hardware Universe.

- 5. Insert the new drive:
 - a. With the cam handle in the open position, use both hands to insert the new drive.
 - b. Push until the drive stops.
 - c. Close the cam handle so that the drive is fully seated into the mid plane and the handle clicks into place. Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.
- 6. Verify that the drive's activity LED (green) is illuminated.

When the drive's activity LED is solid, it means that the drive has power. When the drive's activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

7. To add another drive, repeat Steps 4 through 6.

The new drives are not recognized until they are assigned to a node. You can assign the new drives

manually, or you can wait for ONTAP to automatically assign the new drives if your node follows the rules for drive auto-assignment.

8. After the new drives have all been recognized, verify that they have been added and their ownership is specified correctly.

Steps to confirm installation

1. Display the list of disks:

storage aggregate show-spare-disks

You should see the new drives, owned by the correct node.

2. Optionally (for ONTAP 9.3 and earlier only), zero the newly added drives:

storage disk zerospares

Drives that have been used previously in an ONTAP local tier must be zeroed before they can be added to another local tier. In ONTAP 9.3 and earlier, zeroing can take hours to complete, depending on the size of the non-zeroed drives in the node. Zeroing the drives now can prevent delays in case you need to quickly increase the size of an local tier. This is not an issue in ONTAP 9.4 or later where drives are zeroed using *fast zeroing* which takes only seconds.

Results

The new drives are ready. You can add them to a local tier, place them onto the list of hot spares, or add them when you create a new local tier.

Related information

• storage aggregate show-spare-disks

Correct misaligned ONTAP spare partitions

When you add partitioned disks to a local tier, you must leave a disk with both the root and data partition available as a spare for every node. If you do not and your node experiences a disruption, ONTAP cannot dump the core to the spare data partition.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

Before you begin

You must have both a spare data partition and a spare root partition on the same type of disk owned by the same node.

Steps

1. Using the CLI, display the spare partitions for the node:

```
storage aggregate show-spare-disks -original-owner node_name
```

Note which disk has a spare data partition (spare_data) and which disk has a spare root partition (spare_root). The spare partition will show a non-zero value under the Local Data Usable or Local Root Usable column.

2. Replace the disk with a spare data partition with the disk with the spare root partition:

storage disk replace -disk spare data -replacement spare root -action start

You can copy the data in either direction; however, copying the root partition takes less time to complete.

3. Monitor the progress of the disk replacement:

storage aggregate show-status -aggregate aggr name

4. After the replacement operation is complete, display the spares again to confirm that you have a full spare disk:

```
storage aggregate show-spare-disks -original-owner node name
```

You should see a spare disk with usable space under both "Local Data Usable" and Local Root Usable.

Example

You display your spare partitions for node c1-01 and see that your spare partitions are not aligned:

```
c1::> storage aggregate show-spare-disks -original-owner c1-01
Original Owner: c1-01
 Pool0
 Shared HDD Spares
                          Local
                                  Local
                                  Root Physical
                           Data
 Disk
        Type
              RPM Checksum Usable
                                 Usable
                                            Size
                 _____ _ ____
 _____
                         753.8GB
 1.0.1
        BSAS 7200 block
                                    0B
                                        828.0GB
 1.0.10 BSAS 7200 block
                              OB 73.89GB 828.0GB
```

You start the disk replacement job:

cl::> storage disk replace -disk 1.0.1 -replacement 1.0.10 -action start

While you are waiting for the replacement operation to finish, you display the progress of the operation:

```
c1::> storage aggregate show-status -aggregate aggr0 1
Owner Node: c1-01
 Aggregate: aggr0 1 (online, raid dp) (block checksums)
  Plex: /aggr0 1/plex0 (online, normal, active, pool0)
   RAID Group /aggr0 1/plex0/rg0 (normal, block checksums)
                                   Usable Physical
 Position Disk
                 Pool Type
                             RPM
                                     Size
                                              Size Status
                                               ___ ____
 shared
         1.0.1
                  \cap
                      BSAS
                           7200 73.89GB 828.0GB (replacing, copy in
progress)
 shared
        1.0.10
                  0
                      BSAS 7200 73.89GB 828.0GB (copy 63% completed)
 shared 1.0.0
                  0
                      BSAS 7200 73.89GB 828.0GB (normal)
 shared 1.0.11
                      BSAS
                  0
                           7200 73.89GB 828.0GB (normal)
 shared 1.0.6
                  0
                      BSAS
                           7200 73.89GB 828.0GB (normal)
 shared 1.0.5
                           7200 73.89GB 828.0GB (normal)
                  0
                      BSAS
```

After the replacement operation is complete, confirm that you have a full spare disk:

```
ie2220::> storage aggregate show-spare-disks -original-owner c1-01
Original Owner: c1-01
Pool0
Shared HDD Spares
Local Local
Data Root Physical
Disk Type RPM Checksum Usable Usable Size
------
1.0.1 BSAS 7200 block 753.8GB 73.89GB 828.0GB
```

Related information

• storage aggregate show

Manage disks

How ONTAP hot spare disks work

A hot spare disk is a disk that is assigned to a storage system and is ready for use, but is not in use by a RAID group and does not hold any data.

If a disk failure occurs within a RAID group, the hot spare disk is automatically assigned to the RAID group to replace the failed disks. The data of the failed disk is reconstructed on the hot spare replacement disk in the background from the RAID parity disk. The reconstruction activity is logged in the /etc/message file and an AutoSupport message is sent.

If the available hot spare disk is not the same size as the failed disk, a disk of the next larger size is chosen

and then downsized to match the size of the disk that it is replacing.

Spare requirements for multi-disk carrier disk

Maintaining the proper number of spares for disks in multi-disk carriers is critical for optimizing storage redundancy and minimizing the amount of time that ONTAP must spend copying disks to achieve an optimal disk layout.

You must maintain a minimum of two hot spares for multi-disk carrier disks at all times. To support the use of the Maintenance Center and to avoid issues caused by multiple concurrent disk failures, you should maintain at least four hot spares for steady state operation, and replace failed disks promptly.

If two disks fail at the same time with only two available hot spares, ONTAP might not be able to swap the contents of both the failed disk and its carrier mate to the spare disks. This scenario is called a stalemate. If this happens, you are notified through EMS messages and AutoSupport messages. When the replacement carriers become available, you must follow the instructions that are provided by the EMS messages. For more information, see Knowledge Base article RAID Layout Cannot Be Autocorrected - AutoSupport message

How low spare warnings can help you manage your ONTAP spare disks

By default, warnings are issued to the console and logs if you have fewer than one hot spare drive that matches the attributes of each drive in your storage system.

You can change the threshold value for these warning messages to ensure that your system adheres to best practices.

About this task

You should set the "min_spare_count" RAID option to "2" to ensure that you always have the minimum recommended number of spare disks.

Step

1. Set the option to "2":

storage raid-options modify -node nodename -name min_spare_count -value 2

Additional ONTAP root-data partitioning management options

A root-data partitioning option is available from the Boot Menu that provides additional management features for disks that are configured for root-data partitioning.

The following management features are available under the Boot Menu Option 9.

Unpartition all disks and remove their ownership information

This option is useful if your system is configured for root-data partitioning and you need to reinitialize it with a different configuration.

Clean configuration and initialize node with partitioned disks

This option is useful for the following:

• Your system is not configured for root-data partitioning and you would like to configure it for root-data

partitioning

- Your system is incorrectly configured for root-data partitioning and you need to correct it
- You have an AFF platform or a FAS platform with only SSDs attached that is configured for the previous version of root-data partitioning and you want to upgrade it to the newer version of root-data partitioning to gain increased storage efficiency
- Clean configuration and initialize node with whole disks

This option is useful if you need to:

- Unpartition existing partitions
- Remove local disk ownership
- Reinitialize your system with whole disks using RAID-DP

Learn when to update the ONTAP Disk Qualification Package

The Disk Qualification Package (DQP) adds full support for newly qualified drives. Before you update drive firmware or add new drive types or sizes to a cluster, you must update the DQP. A best practice is to also update the DQP regularly; for example, every quarter or semi-annually.

You need to download and install the DQP in the following situations:

• Whenever you add a new drive type or size to the node

For example, if you already have 1-TB drives and add 2-TB drives, you need to check for the latest DQP update.

- Whenever you update the disk firmware
- Whenever newer disk firmware or DQP files are available
- Whenever you upgrade to a new version of ONTAP.

The DQP is not updated as part of an ONTAP upgrade.

Related information

NetApp Downloads: Disk Qualification Package

NetApp Downloads: Disk Drive Firmware

Disk and partition ownership

Manage the ownership of ONTAP disks and partitions

You can manage the ownership of disks and partitions.

You can perform the following tasks:

Display disk and partition ownership

You can view disk ownership to determine which node controls the storage. You can also view the partition ownership on systems that use shared disks.

Change settings for automatic assignment of disk ownership

You can select a non-default policy for automatically assigning disk ownership or disable automatic assignment of disk ownership.

Manually assign ownership of unpartitioned disks

If your cluster is not configured to use automatic disk ownership assignment, you must assign ownership manually.

Manually assign ownership of partitioned disks

You can set the ownership of the container disk or the partitions manually or by using auto-assignment—just as you do for unpartitioned disks.

Remove a failed disk

A disk that has failed completely is no longer considered by ONTAP to be a usable disk, and you can immediately disconnect the disk from the shelf.

Remove ownership from a disk

ONTAP writes disk ownership information to the disk. Before you remove a spare disk or its shelf from a node, you should remove its ownership information so that it can be properly integrated into another node.

Learn about automatic assignment of ONTAP disk ownership

The automatic assignment of unowned disks is enabled by default. Automatic disk ownership assignments occur 10 minutes after HA pair initialization and every five minutes during normal system operation.

When you add a new disk to an HA pair, for example, when replacing a failed disk, responding to a "low spares" message, or adding capacity, the default auto-assignment policy assigns ownership of the disk to a node as a spare.

The default auto-assignment policy is based on platform-specific characteristics, or the DS460C shelf if your HA pair has only these shelves, and it uses one of the following methods (policies) to assign disk ownership:

Assignment method	Effect on node assignments	Platform configurations that default to the assignment method
bay	Even-numbered bays are assigned to node A and odd-numbered bays to node B.	Entry-level systems in an HA pair configuration with a single, shared shelf.
shelf	All disks in the shelf are assigned to node A.	Entry-level systems in an HA pair configuration with one stack of two or more shelves, and MetroCluster configurations with one stack per node, two or more shelves.

split shelf This policy falls under the "default" value for the -autoassign -policy parameter of the storage disk option command for applicable platform and shelf configurations.	Disks on the left side of the shelf are assigned to node A and on the right side to Node B. Partial shelves on HA pairs are shipped from the factory with disks populated from the shelf edge toward the center.	Most AFF platforms and some MetroCluster configurations.
stack	All disks in the stack are assigned to node A.	Stand-alone entry-level systems and all other configurations.
half-drawer This policy falls under the "default" value for the -autoassign -policy parameter of the storage disk option command for applicable platform and shelf configurations.	All drives in the left half of a DS460C drawer (drive bays 0 to 5) are assigned to node A; all drives in the right half of a drawer (drive bays 6 to 11) are assigned to node B. When initializing an HA pair with only DS460C shelves, automatic assignment of disk ownership is not supported. You must manually assign ownership for drives containing root/container drives that have the root partition by conforming to the half-drawer policy.	HA pairs with only DS460C shelves, after HA pair initialization (boot up). After an HA pair boots up, automatic assignment of disk ownership is automatically enabled and uses the half-drawer policy to assign ownership to the remaining drives (other than the root drives/container drives that have the root partition) and any drives added in the future. If your HA pair has DS460C shelves in addition to other shelf models, the half-drawer policy is not used. The default policy used is dictated by platform-specific characteristics.

Auto-assignment settings and modifications:

- You can display the current auto-assignment settings (on/off) with the storage disk option show command.
- You can disable automatic assignment by using the storage disk option modify command.
- If the default auto-assignment policy is not desirable in your environment, you can specify (change) the bay, shelf, or stack assignment method using the -autoassign-policy parameter in the storage disk option modify command.

Learn how to Change settings for automatic assignment of disk ownership.



The half-drawer and split-shelf default auto-assignment policies are unique because they cannot be set by users like the bay, shelf, and stack policies can.

In Advanced Drive Partitioning (ADP) systems, to make auto-assign work on half-populated shelves, drives must be installed in the correct shelf bays based on what type of shelf you have:

- If your shelf is not a DS460C shelf, install drives equally on the far left side and far right side moving toward the middle. For example, six drives in bays 0-5 and six drives in bays 18-23 of a DS224C shelf.
- If your shelf is a DS460C shelf, install drives in the front row (drive bays 0, 3, 6, and 9) of each drawer. For the remaining drives, evenly distribute them across each drawer by filling drawer rows from front to back. If you don't have enough drives to fill rows, then install them in pairs so that drives occupy the left and right side of a drawer evenly.

Installing drives in the front row of each drawer allows for proper air flow and prevents overheating.



If drives are not installed in the correct shelf bays on half-populated shelves, when a container drive fails and is replaced, ONTAP does not auto-assign ownership. In this case, assignment of the new container drive needs to be done manually. After you have assigned ownership for the container drive, ONTAP automatically handles any drive partitioning and partitioning assignments that are required.

In some situations where auto-assignment will not work, you need to manually assign disk ownership using the storage disk assign command:

- If you disable auto-assignment, new disks are not available as spares until they are manually assigned to a node.
- If you want disks to be auto-assigned and you have multiple stacks or shelves that must have different ownership, one disk must have been manually assigned on each stack or shelf so that automatic ownership assignment works on each stack or shelf.
- If auto-assignment is enabled and you manually assign a single drive to a node that isn't specified in the active policy, auto-assignment stops working and an EMS message is displayed.

Learn how to Manually assign disk ownership of unpartitioned disks.

Learn how to Manually assign disk ownership of partitioned disks.

Display ONTAP disk and partition ownership

You can view disk ownership to determine which node controls the storage. You can also view the partition ownership on systems that use shared disks.

Steps

1. Display the ownership of physical disks:

storage disk show -ownership

cluster::> storage disk show -ownership Disk Aggregate Home Owner DR Home Home ID Owner ID DR Home ID Reserver Pool ----- ------ ------ ------_____ _ _____ ____ 1.0.0 aggr0 2 node2 node2 - 2014941509 2014941509 -2014941509 Pool0 1.0.1 aggr0 2 node2 node2 - 2014941509 2014941509 -2014941509 Pool0 1.0.2 aggr0 1 node1 node1 _ 2014941219 2014941219 -2014941219 Pool0 1.0.3 nodel nodel 2014941219 2014941219 _ _ 2014941219 Pool0

2. If you have a system that uses shared disks, you can display the partition ownership:

storage disk show -partition-ownership

cluster::> storage	disk show -	partition-ow	nership		
		Root		Data	
Container Contain	er				
Disk Aggregate	Root Owner	Owner ID	Data Owner	Owner ID	Owner
Owner ID					
1.0.0 -	node1	1886742616	node1	1886742616	node1
1886742616					
1.0.1 -	nodel	1886742616	nodel	1886742616	node1
1886742616					
1.0.2 -	node2	1886742657	node2	1886742657	node2
1886742657					
1.0.3 -	node2	1886742657	node2	1886742657	node2
1886742657					

Change settings for automatic assignment of ONTAP disk ownership

You can use the storage disk option modify command to select a non-default policy for automatically assigning disk ownership or to disable automatic assignment of disk ownership.

Learn about automatic assignment of disk ownership.

About this task

If you have an HA pair with only DS460C shelves, the default auto-assignment policy is half-drawer. You cannot change to a non-default policy (bay, shelf, stack).

Steps

- 1. Modify automatic disk assignment:
 - a. If you want to select a non-default policy, enter:

```
storage disk option modify -autoassign-policy autoassign_policy -node
node_name
```

- Use stack as the *autoassign_policy* to configure automatic ownership at the stack or loop level.
- Use shelf as the *autoassign policy* to configure automatic ownership at the shelf level.
- Use bay as the *autoassign* policy to configure automatic ownership at the bay level.
- b. If you want to disable automatic disk ownership assignment, enter:

storage disk option modify -autoassign off -node node_name

2. Verify the automatic assignment settings for the disks:

storage disk option show

Auto Assign	Auto Assign Policy
on on	default default
	Auto Assign on on

Manually assign ONTAP disk ownership of unpartitioned disks

If your HA pair is not configured to use automatic disk ownership assignment, you must manually assign ownership. If you are initializing an HA pair that has only DS460C shelves, you must manually assign ownership for the root drives.

About this task

- If you are manually assigning ownership in an HA pair that is not being initialized and does not have only DS460C shelves, use option 1.
- If you are initializing an HA pair that has only DS460C shelves, use option 2 to manually assign ownership for the root drives.

Option 1: Most HA pairs

For an HA pair that is not being initialized and does not have only DS460C shelves, use this procedure to manually assigning ownership.

About this task

- The disks you are assigning ownership for must be in a shelf that is physically cabled to the node you are assigning ownership to.
- If you are using disks in a local tier (aggregate):
 - Disks must be owned by a node before they can be used in a local tier (aggregate).
 - You cannot reassign ownership of a disk that is in use in a local tier (aggregate).

Steps

1. Use the CLI to display all unowned disks:

storage disk show -container-type unassigned

2. Assign each disk:

storage disk assign -disk disk name -owner owner name

You can use the wildcard character to assign more than one disk at once. If you are reassigning a spare disk that is already owned by a different node, you must use the "-force" option.

For an HA pair that you are initializing and that only has DS460C shelves, use this procedure to manually assign ownership for the root drives.

About this task

• When you initialize an HA pair that has only DS460C shelves, you must manually assign the root drives to conform to the half-drawer policy.

After HA pair initialization (boot up), automatic assignment of disk ownership is automatically enabled and uses the half-drawer policy to assign ownership to the remaining drives (other than the root drives) and any drives added in the future, such as replacing failed disks, responding to a "low spares" message, or adding capacity.

Learn about the half-drawer policy.

• RAID needs a minimum of 10 drives for each HA pair (5 for each node) for any greater than 8TB NL-SAS drives in a DS460C shelf.

Steps

- 1. If your DS460C shelves are not fully populated, complete the following substeps; otherwise, go to the next step.
 - a. First, install drives in the front row (drive bays 0, 3, 6, and 9) of each drawer.

Installing drives in the front row of each drawer allows for proper air flow and prevents overheating.

b. For the remaining drives, evenly distribute them across each drawer.

Fill drawer rows from front to back. If you don't have enough drives to fill rows, then install them in pairs so that drives occupy the left and right side of a drawer evenly.

The following illustration shows the drive bay numbering and locations in a DS460C drawer.



2. Log into the clustershell using the node-management LIF or cluster-management LIF.

3. Manually assign the root drives in each drawer to conform to the half-drawer policy using the following substeps:

The half-drawer policy has you assign the left half of a drawer's drives (bays 0 to 5) to node A, and the right half of a drawer's drives (bays 6 to 11) to node B.

- a. Display all unowned disks: storage disk show -container-type unassigned
- b. Assign the root disks: storage disk assign -disk disk name -owner owner name

You can use the wildcard character to assign more than one disk at a time.

Learn more about storage disk in the ONTAP command reference.

Manually assign ownership of ONTAP partitioned disks

You can manually assign the ownership of the container disk or the partitions on Advanced Drive Partitioning (ADP) systems. If you are initializing an HA pair that only has DS460C shelves, you must manually assign ownership for the container drives that will include root partitions.

About this task

• The type of storage system you have determines which method of ADP is supported, root-data (RD) or root-data-data (RD2).

FAS storage systems use RD and AFF storage systems use RD2.

- If you are manually assigning ownership in an HA pair that is not being initialized and does not have only DS460C shelves, use option 1 to manually assign disks with root-data (RD) partitioning or use option 2 to manually assign disks with root-data-data (RD2) partitioning.
- If you are initializing an HA pair that has only DS460C shelves, use option 3 to manually assign ownership for the container drives that have the root partition.

For root-data partitioning, there are three owned entities (the container disk and the two partitions) collectively owned by the HA pair.

About this task

- The container disk and the two partitions do not all need to be owned by the same node in the HA pair as long as they are all owned by one of the nodes in the HA pair. However, when you use a partition in a local tier, it must be owned by the same node that owns the local tier.
- If a container disk fails in a half-populated shelf and is replaced, you might need to manually assign disk ownership because ONTAP does not always auto-assign ownership in this case.
- After the container disk is assigned, ONTAP's software automatically handles any partitioning and partition assignments that are required.

Steps

1. Use the CLI to display the current ownership for the partitioned disk:

storage disk show -disk disk_name -partition-ownership

2. Set the CLI privilege level to advanced:

set -privilege advanced

3. Enter the appropriate command, depending on which ownership entity you want to assign ownership for:

If any of the ownership entities are already owned, then you must include the -force option.

If you want to assign ownership for the	Use this command
Container disk	storage disk assign -disk <i>disk_name</i> -owner <i>owner_name</i>
Data partition	storage disk assign -disk <i>disk_name</i> -owner <i>owner_name</i> -data true
Root partition	storage disk assign -disk <i>disk_name</i> -owner <i>owner_name</i> -root true

For root-data-data partitioning, there are four owned entities (the container disk and the three partitions) collectively owned by the HA pair. Root-data-data partitioning creates one small partition as the root partition and two larger, equally sized partitions for data.

About this task

- Parameters must be used with the disk assign command to assign the proper partition of a rootdata-data partitioned disk. You cannot use these parameters with disks that are part of a storage pool. The default value is false.
 - The -data1 true parameter assigns the data1 partition of a root-data1-data2 partitioned disk.
 - The -data2 true parameter assigns the data2 partition of a root-data1-data2 partitioned disk.
- If a container disk fails in a half-populated shelf and is replaced, you might need to manually assign disk ownership because ONTAP does not always auto-assign ownership in this case.
- After the container disk is assigned, ONTAP's software automatically handles any partitioning and partition assignments that are required.

Steps

1. Use the CLI to display the current ownership for the partitioned disk:

storage disk show -disk disk name -partition-ownership

2. Set the CLI privilege level to advanced:

set -privilege advanced

3. Enter the appropriate command, depending on which ownership entity you want to assign ownership for:

If any of the ownership entities are already owned, then you must include the -1010e optic
--

If you want to assign ownership for the	Use this command
Container disk	storage disk assign -disk <i>disk_name</i> -owner <i>owner_name</i>
Data1 partition	storage disk assign -disk <i>disk_name</i> -owner <i>owner_name</i> -data1 true
Data2 partition	storage disk assign -disk <i>disk_name</i> -owner <i>owner_name</i> -data2 true
Root partition	storage disk assign -disk <i>disk_name</i> -owner <i>owner_name</i> -root true

If you are initializing an HA pair that has only DS460C shelves, you must manually assign ownership for the container drives that have the root partition by conforming to the half-drawer policy.

About this task

• When you initialize an HA pair that has only DS460C shelves, the ADP boot menu options 9a and 9b do not support automatic drive ownership assignment. You must manually assign the container drives that have the root partition by conforming to the half-drawer policy.

After HA pair initialization (boot up), automatic assignment of disk ownership is automatically enabled and uses the half-drawer policy to assign ownership to the remaining drives (other than the container drives that have the root partition) and any drives added in the future, such as replacing failed drives, responding to a "low spares" message, or adding capacity.

• Learn about the half-drawer policy.

Steps

- 1. If your DS460C shelves are not fully populated, complete the following substeps; otherwise, go to the next step.
 - a. First, install drives in the front row (drive bays 0, 3, 6, and 9) of each drawer.

Installing drives in the front row of each drawer allows for proper air flow and prevents overheating.

b. For the remaining drives, evenly distribute them across each drawer.

Fill drawer rows from front to back. If you don't have enough drives to fill rows, then install them in pairs so that drives occupy the left and right side of a drawer evenly.

The following illustration shows the drive bay numbering and locations in a DS460C drawer.



- 2. Log into the clustershell using the node-management LIF or cluster-management LIF.
- 3. For each drawer, manually assign the container drives that have the root partition by conforming to the half-drawer policy using the following substeps:

The half-drawer policy has you assign the left half of a drawer's drives (bays 0 to 5) to node A, and the right half of a drawer's drives (bays 6 to 11) to node B.

- a. Display all unowned disks: storage disk show -container-type unassigned
- b. Assign the container drives that have the root partition: storage disk assign -disk disk_name -owner owner_name

You can use the wildcard character to assign more than one drive at a time.

Set up an active-passive configuration on ONTAP nodes using root-data partitioning

When an HA pair is configured to use root-data partitioning by the factory, ownership of the data partitions is split between both nodes in the pair for use in an active-active configuration. If you want to use the HA pair in an active-passive configuration, you must update partition ownership before creating your data local tier.

Before you begin

- You should have decided which node will be the active node and which node will be the passive node.
- Storage failover must be configured on the HA pair.

About this task

This task is performed on two nodes: Node A and Node B.

This procedure is designed for nodes for which no data local tier has been created from the partitioned disks.

Learn about advanced disk partitioning.

Steps

All commands are inputted at the cluster shell.

1. View the current ownership of the data partitions:

storage aggregate show-spare-disks

The output shows that half of the data partitions are owned by one node and half are owned by the other node. All of the data partitions should be spare.

```
cluster1::> storage aggregate show-spare-disks
Original Owner: cluster1-01
Pool0
Partitioned Spares
Local
Local
Data
Root Physical
Disk Type RPM Checksum Usable
```

Usable Size				
1.0.0	BSAS	7200	block	753.8GB
0B 828.0GB		7200	h l a c h	762 000
1.0.1 73.89GB 828.0GB	BSAS	1200	DIOCK	/33.8GB
1.0.5	BSAS	7200	block	753.8GB
0B 828.0GB	DCJC	7200	block	753 8CD
0B 828.0GB	DOAD	1200	DIOCK	/ JJ . 0GD
1.0.10	BSAS	7200	block	753.8GB
OB 828.0GB 1.0.11	BSAS	7200	block	753.8GB
0B 828.0GB	20110	0 0		,
Original Owner: cluster1-02				
Pool0				
Partitioned Spares				
Local				Local
TOOGT				
				Data
Root Physical	Turne	DDM	Chookoum	Data
Root Physical Disk Usable Size	Туре	RPM	Checksum	Data Usable
Root Physical Disk Usable Size	Туре	RPM	Checksum	Data Usable
Root Physical Disk Usable Size 	Type 	RPM	Checksum block	Data Usable
Root Physical Disk Usable Size 1.0.2 OB 828.0GB	Type BSAS	RPM 7200	Checksum block	Data Usable 753.8GB
Root Physical Disk Usable Size 1.0.2 OB 828.0GB 1.0.3 OB 828.0GB	Type 	RPM 7200 7200	Checksum block block	Data Usable 753.8GB 753.8GB
Root Physical Disk Usable Size 	Type 	RPM 7200 7200 7200	Checksum block block block	Data Usable 753.8GB 753.8GB 753.8GB
Root Physical Disk Usable Size 1.0.2 OB 828.0GB 1.0.3 OB 828.0GB 1.0.4 OB 828.0GB	Type 	RPM 7200 7200 7200	Checksum block block block	Data Usable 753.8GB 753.8GB 753.8GB
Root Physical Disk Usable Size 1.0.2 OB 828.0GB 1.0.3 OB 828.0GB 1.0.4 OB 828.0GB 1.0.7 OB 828.0GB	Type BSAS BSAS BSAS BSAS BSAS	RPM 7200 7200 7200 7200	Checksum block block block block	Data Usable 753.8GB 753.8GB 753.8GB 753.8GB
Root Physical Disk Usable Size 1.0.2 OB 828.0GB 1.0.3 OB 828.0GB 1.0.4 OB 828.0GB 1.0.7 OB 828.0GB 1.0.7 OB 828.0GB 1.0.8	Type BSAS BSAS BSAS BSAS BSAS	RPM 7200 7200 7200 7200 7200	Checksum block block block block block	Data Usable 753.8GB 753.8GB 753.8GB 753.8GB 753.8GB
Root Physical Disk Usable Size 1.0.2 OB 828.0GB 1.0.3 OB 828.0GB 1.0.4 OB 828.0GB 1.0.7 OB 828.0GB 1.0.7 OB 828.0GB 1.0.8 73.89GB 828.0GB	Type BSAS BSAS BSAS BSAS BSAS	RPM 7200 7200 7200 7200 7200	Checksum block block block block block	Data Usable 753.8GB 753.8GB 753.8GB 753.8GB 753.8GB
Root Physical Disk Usable Size 1.0.2 OB 828.0GB 1.0.3 OB 828.0GB 1.0.4 OB 828.0GB 1.0.4 OB 828.0GB 1.0.7 OB 828.0GB 1.0.8 73.89GB 828.0GB 1.0.9 OB 828.0GB	Type BSAS BSAS BSAS BSAS BSAS BSAS	RPM 7200 7200 7200 7200 7200 7200	Checksum block block block block block block	Data Usable 753.8GB 753.8GB 753.8GB 753.8GB 753.8GB 753.8GB
Root Physical Disk Usable Size 1.0.2 OB 828.0GB 1.0.3 OB 828.0GB 1.0.4 OB 828.0GB 1.0.7 OB 828.0GB 1.0.7 OB 828.0GB 1.0.8 73.89GB 828.0GB 1.0.9 OB 828.0GB 1.0.9	Type BSAS BSAS BSAS BSAS BSAS BSAS	RPM 7200 7200 7200 7200 7200 7200	Checksum block block block block block block	Data Usable 753.8GB 753.8GB 753.8GB 753.8GB 753.8GB 753.8GB

2. Enter the advanced privilege level:

set advanced

3. For each data partition owned by the node that will be the passive node, assign it to the active node:

storage disk assign -force -data true -owner active_node_name -disk disk_name

You do not need to include the partition as part of the disk name.

You would enter a command similar to the following example for each data partition you need to reassign:

storage disk assign -force -data true -owner cluster1-01 -disk 1.0.3

4. Confirm that all of the partitions are assigned to the active node.

```
cluster1::*> storage aggregate show-spare-disks
Original Owner: cluster1-01
PoolO
 Partitioned Spares
                                               Local
Local
                                                Data
Root Physical
                      Type RPM Checksum Usable
Disk
Usable Size
 _____ ____
_____ ____
1.0.0
                      BSAS 7200 block
                                             753.8GB
0B 828.0GB
                      BSAS 7200 block
1.0.1
                                             753.8GB
73.89GB 828.0GB
1.0.2
                      BSAS 7200 block
                                             753.8GB
0B 828.0GB
1.0.3
                      BSAS 7200 block
                                             753.8GB
0B 828.0GB
                      BSAS 7200 block
1.0.4
                                             753.8GB
0B 828.0GB
1.0.5
                      BSAS 7200 block
                                             753.8GB
0B 828.0GB
1.0.6
                            7200 block
                                             753.8GB
                      BSAS
0B 828.0GB
1.0.7
                            7200 block
                      BSAS
                                             753.8GB
0B 828.0GB
1.0.8
                            7200 block
                                             753.8GB
                      BSAS
0B 828.0GB
1.0.9
                      BSAS 7200 block
                                             753.8GB
0B 828.0GB
1.0.10
                      BSAS 7200 block
                                             753.8GB
0B 828.0GB
1.0.11
                             7200 block
                                             753.8GB
                      BSAS
0B 828.0GB
```

Original Owner: cluster1-02 Pool0 Partitioned Spares Local Local Data Root Physical Disk Type RPM Checksum Usable Usable Size _____ _____ 1.0.8 BSAS 7200 block 0B 73.89GB 828.0GB 13 entries were displayed.

Note that cluster1-02 still owns a spare root partition.

5. Return to administrative privilege:

set admin

6. Create your data local tier, leaving at least one data partition as spare:

```
storage aggregate create new_aggr_name -diskcount number_of_partitions -node
active_node_name
```

The data local tier is created and is owned by the active node.

Related information

- storage aggregate create
- storage aggregate show

Set up an active-passive configuration on ONTAP nodes using root-data-data partitioning

When an HA pair is configured to use root-data-data partitioning by the factory, ownership of the data partitions is split between both nodes in the pair for use in an active-active configuration. If you want to use the HA pair in an active-passive configuration, you must update partition ownership before creating your data local tier.

Before you begin

- You should have decided which node will be the active node and which node will be the passive node.
- Storage failover must be configured on the HA pair.

About this task

This task is performed on two nodes: Node A and Node B.

This procedure is designed for nodes for which no data local tier has been created from the partitioned disks.

Learn about advanced disk partitioning.

Steps

All commands are input at the cluster shell.

1. View the current ownership of the data partitions:

```
storage aggregate show-spare-disks -original-owner passive_node_name -fields
local-usable-data1-size, local-usable-data2-size
```

The output shows that half of the data partitions are owned by one node and half are owned by the other node. All of the data partitions should be spare.

2. Enter the advanced privilege level:

set advanced

3. For each data1 partition owned by the node that will be the passive node, assign it to the active node:

storage disk assign -force -data1 -owner active_node_name -disk disk_name

You do not need to include the partition as part of the disk name

4. For each data2 partition owned by the node that will be the passive node, assign it to the active node:

storage disk assign -force -data2 -owner active_node_name -disk disk_name

You do not need to include the partition as part of the disk name

5. Confirm that all of the partitions are assigned to the active node:

storage aggregate show-spare-disks

```
cluster1::*> storage aggregate show-spare-disks
Original Owner: cluster1-01
PoolO
 Partitioned Spares
                                              Local
Local
                                               Data
Root Physical
Disk
                      Type RPM Checksum Usable
Usable
       Size
 _____ _____
_____ _ ___
                      BSAS 7200 block 753.8GB
1.0.0
0B 828.0GB
1.0.1
                      BSAS 7200 block
                                            753.8GB
73.89GB 828.0GB
```

1.0.2 OB 828 OCB	BSAS	7200	block	753.8GB
1.0.3	BSAS	7200	block	753.8GB
0B 828.0GB				
1.0.4	BSAS	7200	block	753.8GB
1.0.5	BSAS	7200	block	753.8GB
0B 828.0GB				
1.0.6	BSAS	7200	block	753.8GB
0B 828.0GB	DGJG	7000	l- l l-	752 000
1.U./ 0B 828.0GB	BSAS	1200	DIOCK	/53.8GB
1.0.8	BSAS	7200	block	753.8GB
0B 828.0GB				
1.0.9	BSAS	7200	block	753.8GB
1.0.10	BSAS	7200	block	753.8GB
0B 828.0GB	20110	0 0	220011	
1.0.11	BSAS	7200	block	753.8GB
0B 828.0GB				
Original Owner: cluster1-02				
Pool0				
Partitioned Spares				
Local				LOCAL
				Data
Root Physical				
Disk Usabla Siza	Туре	RPM	Checksum	Usable
1.0.8	BSAS	7200	block	OB
13.89GB 828.UGB				
is entired were displayed.				

Note that cluster1-02 still owns a spare root partition.

6. Return to administrative privilege:

set admin

7. Create your data aggregate, leaving at least one data partition as spare:

storage aggregate create new_aggr_name -diskcount number_of_partitions -node
active_node_name

The data aggregate is created and is owned by the active node.

8. Alternatively, you can use ONTAP's recommended local tier layout which includes best practices for RAID group layout and spare counts:

storage aggregate auto-provision

Related information

- storage aggregate auto-provision
- storage aggregate create
- storage aggregate show

Remove ONTAP ownership from a disk

ONTAP writes disk ownership information to the disk. Before you remove a spare disk or its shelf from a node, you should remove its ownership information so that it can be properly integrated into another node.



If the disk is partitioned for root-data partitioning and you are running ONTAP 9.10.1 or later, contact NetApp Technical Support for assistance in removing ownership. For more information see the Knowledge Base article: Failed to remove the owner of disk.

Before you begin

The disk you want to remove ownership from must meet the following requirements:

• It must be a spare disk.

You cannot remove ownership from a disk that is being used in an local tier.

- It cannot be in the maintenance center.
- It cannot be undergoing sanitization.
- · It cannot have failed.

It is not necessary to remove ownership from a failed disk.

About this task

If you have automatic disk assignment enabled, ONTAP could automatically reassign ownership before you remove the disk from the node. For this reason, you disable the automatic ownership assignment until the disk is removed, and then you re-enable it.

Steps

1. If disk ownership automatic assignment is on, use the CLI to turn it off:

storage disk option modify -node node_name -autoassign off

- 2. If needed, repeat the previous step for the node's HA partner.
- 3. Remove the software ownership information from the disk:

storage disk removeowner disk_name

To remove ownership information from multiple disks, use a comma-separated list.

Example:

storage disk removeowner sys1:0a.23, sys1:0a.24, sys1:0a.25

4. If the disk is partitioned for root-data partitioning and you are running ONTAP 9.9.1 or earlier, remove ownership from the partitions:

```
storage disk removeowner -disk disk_name -root true storage disk removeowner -disk disk_name -data true
```

Both partitions are no longer owned by any node.

5. If you previously turned off automatic assignment of disk ownership, turn it on after the disk has been removed or reassigned:

storage disk option modify -node node_name -autoassign on

6. If needed, repeat the previous step for the node's HA partner.

Remove a failed ONTAP disk

A disk that has completely failed is no longer counted by ONTAP as a usable disk, and you can immediately disconnect the disk from the disk shelf. However, you should leave a partially failed disk connected long enough for the Rapid RAID Recovery process to complete.

About this task

If you are removing a disk because it has failed or because it is producing excessive error messages, you should not use the disk again in this or any other storage system.

Steps

1. Use the CLI to find the disk ID of the failed disk:

storage disk show -broken

If the disk does not appear in the list of failed disks, it might have partially failed, with a Rapid RAID Recovery in process. In this case, you should wait until the disk is present in the list of failed disks (which means that the Rapid RAID Recovery process is complete) before removing the disk.

2. Determine the physical location of the disk you want to remove:

storage disk set-led -action on -disk disk name 2

The fault LED on the face of the disk is lit.

3. Remove the disk from the disk shelf, following the instructions in the hardware guide for your disk shelf model.

Disk sanitization

Learn about ONTAP disk sanitization

Disk sanitization is the process of physically obliterating data by overwriting disks or SSDs with specified byte patterns or random data so that recovery of the original data becomes impossible. Using the sanitization process ensures that no one can recover the data on the disks.

This functionality is available through the nodeshell in all ONTAP 9 releases, and starting with ONTAP 9.6 in maintenance mode.

The disk sanitization process uses three successive default or user-specified byte overwrite patterns for up to seven cycles per operation. The random overwrite pattern is repeated for each cycle.

Depending on the disk capacity, the patterns, and the number of cycles, the process can take several hours. Sanitization runs in the background. You can start, stop, and display the status of the sanitization process. The sanitization process contains two phases: the "Formatting phase" and the "Pattern overwrite phase".

Formatting phase

The operation performed for the formatting phase depends on the class of disk being sanitized, as shown in the following table:

Disk class	Formatting phase operation
Capacity HDDs	Skipped
Performance HDDs	SCSI format operation
SSDs	SCSI sanitize operation

Pattern overwrite phase

The specified overwrite patterns are repeated for the specified number of cycles.

When the sanitization process is complete, the specified disks are in a sanitized state. They are not returned to spare status automatically. You must return the sanitized disks to the spare pool before the newly sanitized disks are available to be added to another local tier.

Learn about when ONTAP disk sanitization cannot be performed

Disk sanitization cannot be performed under these circumstances.

- It is not supported in takeover mode for systems in an HA pair.
- It cannot be performed on disks that were failed due to readability or writability problems.
- If you are using the random pattern, it cannot be performed on more than 100 disks at one time.
- It is not supported on array LUNs.

What happens if ONTAP disk sanitization is interrupted

If disk sanitization is interrupted by user intervention or an unexpected event such as a power outage, ONTAP takes action to return the disks that were being sanitized to a known state, but you must also take action before the sanitization process can finish.

Disk sanitization is a long-running operation. If the sanitization process is interrupted by power failure, system panic, or manual intervention, the sanitization process must be repeated from the beginning. The disk is not designated as sanitized.

If the formatting phase of disk sanitization is interrupted, ONTAP must recover any disks that were corrupted by the interruption. After a system reboot and once every hour, ONTAP checks for any sanitization target disk that did not complete the formatting phase of its sanitization. If any such disks are found, ONTAP recovers them. The recovery method depends on the type of the disk. After a disk is recovered, you can rerun the sanitization process on that disk; for HDDs, you can use the -s option to specify that the formatting phase is not repeated again.

Tips for creating and backing up ONTAP local tiers containing data to be sanitized

If you are creating or backing up local tiers to contain data that might need to be sanitized, following some simple guidelines will reduce the time it takes to sanitize your data.

• Make sure your local tiers containing sensitive data are not larger than they need to be.

If they are larger than needed, sanitization requires more time, disk space, and bandwidth.

• When you back up local tiers containing sensitive data, avoid backing them up to local tier that also contain large amounts of nonsensitive data.

This reduces the resources required to move nonsensitive data before sanitizing sensitive data.

Sanitize an ONTAP disk

Sanitizing a disk allows you to remove data from a disk or a set of disks on decommissioned or inoperable systems so that the data can never be recovered.

Two methods are available to sanitize disks using the CLI:

Beginning with ONTAP 9.6, you can perform disk sanitization in maintenance mode.

Before you begin

• The disks cannot be self-encrypting disks (SED).

You must use the storage encryption disk sanitize command to sanitize an SED.

Encryption of data at rest

Learn more about storage encryption disk sanitize in the ONTAP command reference.

Steps

- 1. Boot into maintenance mode.
 - a. Exit the current shell by entering halt.

The LOADER prompt is displayed.

b. Enter maintenance mode by entering boot_ontap maint.

After some information is displayed, the maintenance mode prompt is displayed.

2. If the disks you want to sanitize are partitioned, unpartition each disk:



The command to unpartition a disk is only available at the diag level and should be performed only under NetApp Support supervision. It is highly recommended that you contact NetApp Support before you proceed.

You can also refer to the Knowledge Base article How to unpartition a spare drive in ONTAP

```
disk unpartition <disk name>
```

3. Sanitize the specified disks:

```
disk sanitize start [-p <pattern1>|-r [-p <pattern2>|-r [-p <pattern3>|-
r]]] [-c <cycle count>] <disk list>
```



Do not turn off power to the node, disrupt the storage connectivity, or remove target disks while sanitizing. If sanitizing is interrupted during the formatting phase, the formatting phase must be restarted and allowed to finish before the disks are sanitized and ready to be returned to the spare pool. If you need to abort the sanitization process, you can do so by using the disk sanitize abort command. If the specified disks are undergoing the formatting phase of sanitization, the abort does not occur until the phase is complete.

-p < pattern1 > -p < pattern2 > -p < pattern3 > specifies a cycle of one to three user-defined hexbyte overwrite patterns that can be applied in succession to the disks being sanitized. The defaultpattern is three passes, using 0x55 for the first pass, 0xaa for the second pass, and 0x3c for the thirdpass. -r replaces a patterned overwrite with a random overwrite for any or all of the passes.

-c <cycle_count> specifies the number of times that the specified overwrite patterns are applied. The default value is one cycle. The maximum value is seven cycles.

<disk list> specifies a space-separated list of the IDs of the spare disks to be sanitized.

4. If desired, check the status of the disk sanitization process:

disk sanitize status [<disk_list>]

5. After the sanitization process is complete, return the disks to spare status for each disk:

disk sanitize release <disk name>

6. Exit maintenance mode.

After the disk sanitization feature is enabled using nodeshell commands on a node, it cannot be disabled.

Before you begin

• The disks must be spare disks; they must be owned by a node, but not used in a local tier.

If the disks are partitioned, neither partition can be in use in a local tier.

• The disks cannot be self-encrypting disks (SED).

You must use the storage encryption disk sanitize command to sanitize an SED.

Encryption of data at rest

• The disks cannot be part of a storage pool.

Steps

1. If the disks you want to sanitize are partitioned, unpartition each disk:



The command to unpartition a disk is only available at the diag level and should be performed only under NetApp Support supervision. It is highly recommended that you contact NetApp Support before you proceed. You can also refer to the Knowledge Base article How to unpartition a spare drive in ONTAP.

disk unpartition <disk name>

2. Enter the nodeshell for the node that owns the disks you want to sanitize:

system node run -node <node_name>

3. Enable disk sanitization:

options licensed_feature.disk_sanitization.enable on

You are asked to confirm the command because it is irreversible.

4. Switch to the nodeshell advanced privilege level:

priv set advanced

5. Sanitize the specified disks:

```
disk sanitize start [-p <pattern1>|-r [-p <pattern2>|-r [-p <pattern3>|-
r]]] [-c <cycle_count>] <disk_list>
```

Do not turn off power to the node, disrupt the storage connectivity, or remove target disks while sanitizing. If sanitizing is interrupted during the formatting phase, the formatting

phase must be restarted and allowed to finish before the disks are sanitized and ready to be

returned to the spare pool. If you need to abort the sanitization process, you can do so by using the disk sanitize

abort command. If the specified disks are undergoing the formatting phase of sanitization, the

abort does not occur until the phase is complete.

-p <pattern1> -p <pattern2> -p <pattern3> specifies a cycle of one to three user-defined hex byte

overwrite patterns that can be applied in succession to the disks being sanitized. The default pattern is three passes, using 0x55 for the first pass, 0xaa for the second pass, and 0x3c for the third pass.

-r replaces a patterned overwrite with a random overwrite for any or all of the passes.

-c <cycle_count> specifies the number of times that the specified overwrite patterns are applied.

The default value is one cycle. The maximum value is seven cycles.

<disk list> specifies a space-separated list of the IDs of the spare disks to be sanitized.

6. If you want to check the status of the disk sanitization process:

disk sanitize status [<disk_list>]

7. After the sanitization process is complete, return the disks to spare status:

disk sanitize release <disk_name>

8. Return to the nodeshell admin privilege level:

priv set admin

9. Return to the ONTAP CLI:

exit

10. Determine whether all of the disks were returned to spare status:

storage aggregate show-spare-disks

lf	Then
All of the sanitized disks are listed as spares	You are done. The disks are sanitized and in spare status.

Some of the sanitized disks are	Complete the following steps:
not listed as spares	a. Enter advanced privilege mode:
	set -privilege advanced
	b. Assign the unassigned sanitized disks to the appropriate nod for each disk:
	storage disk assign -disk <disk_name> -owner <node_name></node_name></disk_name>
	c. Return the disks to spare status for each disk:
	storage disk unfail -disk <disk_name> -s -q</disk_name>
	d. Return to administrative mode:
	set -privilege admin

Learn more about storage aggregate show-spare-disks in the ONTAP command reference.

Result

The specified disks are sanitized and designated as hot spares. The serial numbers of the sanitized disks are written to /etc/log/sanitized_disks.

The specified disks' sanitization logs, which show what was completed on each disk, are written to /mroot/etc/log/sanitization.log.

Related information

• storage aggregate show

ONTAP commands for managing disks

You can use the storage disk and storage aggregate commands to manage your disks.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

If you want to	Use this command
Display a list of spare disks, including partitioned disks, by owner	storage aggregate show-spare-disks
Display the disk RAID type, current usage, and RAID group by local tier	storage aggregate show-status
Display the RAID type, current usage, local tier, and RAID group, including spares, for physical disks	storage disk show -raid
--	---
Display a list of failed disks	storage disk show -broken
Display the pre-cluster (nodescope) drive name for a disk	storage disk show -primary-paths (advanced)
Illuminate the LED for a particular disk or shelf	storage disk set-led
Display the checksum type for a specific disk	storage disk show -fields checksum- compatibility
Display the checksum type for all spare disks	storage disk show -fields checksum- compatibility -container-type spare
Display disk connectivity and placement information	storage disk show -fields disk,primary- port,secondary-name,secondary- port,shelf,bay
Display the pre-cluster disk names for specific disks	storage disk show -disk diskname -fields diskpathnames
Display the list of disks in the maintenance center	storage disk show -maintenance
Display SSD wear life	storage disk show -ssd-wear
Unpartition a shared disk	storage disk unpartition (available at diagnostic level)
Zero all non-zeroed disks	storage disk zerospares
Stop an ongoing sanitization process on one or more specified disks	system node run -node nodename -command disk sanitize
Display storage encryption disk information	storage encryption disk show
Retrieve authentication keys from all linked key management servers	security key-manager restore

Related information

• storage aggregate show

ONTAP commands for displaying space usage information

You use the storage aggregate and volume commands to see how space is being used in your local tiers and volumes and their snapshots.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

To display information about	Use this command
Local tier, including details about used and available space percentages, snapshot reserve size, and other space usage information	<pre>storage aggregate show storage aggregate show-space -fields snap-size-total,used-including- snapshot-reserve</pre>
How disks and RAID groups are used in an local tier, and RAID status	storage aggregate show-status
The amount of disk space that would be reclaimed if you deleted a specific snapshot	volume snapshot compute-reclaimable
The amount of space used by a volume	<pre>volume show -fields size,used,available,percent-used volume show-space</pre>
The amount of space used by a volume in the containing local tier	volume show-footprint

Related information

- storage aggregate show
- storage aggregate show-space
- storage aggregate show-status

ONTAP commands for displaying information about storage shelves

You use the storage shelf show command to display configuration and error information for your disk shelves.

If you want to display	Use this command
General information about shelf configuration and hardware status	storage shelf show

If you want to display	Use this command
Detailed information for a specific shelf, including stack ID	storage shelf show -shelf
Unresolved, customer actionable, errors by shelf	storage shelf show -errors
Bay information	storage shelf show -bay
Connectivity information	storage shelf show -connectivity
Cooling information, including temperature sensors and cooling fans	storage shelf show -cooling
Information about I/O modules	storage shelf show -module
Port information	storage shelf show -port
Power information, including PSUs (power supply units), current sensors, and voltage sensors	storage shelf show -power

Related information

ONTAP command reference

Manage RAID configurations

Default RAID policies for ONTAP local tiers

Either RAID-DP or RAID-TEC is the default RAID policy for all new local tiers. The RAID policy determines the parity protection you have in the event of a disk failure.

RAID-DP provides double-parity protection in the event of a single or double disk failure. RAID-DP is the default RAID policy for the following local tier types:

- All Flash local tiers
- Flash Pool local tiers
- Performance hard disk drive (HDD) local tiers

RAID-TEC is supported on all disk types and all platforms, including AFF. Local tiers that contain larger disks have a higher possibility of concurrent disk failures. RAID-TEC helps to mitigate this risk by providing triple-parity protection so that your data can survive up to three simultaneous disk failures. RAID-TEC is the default RAID policy for capacity HDD local tiers with disks that are 6 TB or larger.

Each RAID policy type requires a minimum number of disks:

- RAID-DP: minimum of 5 disks
- RAID-TEC: minimum of 7 disks

ONTAP RAID protection levels for disks

ONTAP supports three levels of RAID protection for local tiers. The level of RAID protection determines the number of parity disks available for data recovery in the event of disk failures.

With RAID protection, if there is a data disk failure in a RAID group, ONTAP can replace the failed disk with a spare disk and use parity data to reconstruct the data of the failed disk.

• RAID4

With RAID4 protection, ONTAP can use one spare disk to replace and reconstruct the data from one failed disk within the RAID group.

• RAID-DP

With RAID-DP protection, ONTAP can use up to two spare disks to replace and reconstruct the data from up to two simultaneously failed disks within the RAID group.

• RAID-TEC

With RAID-TEC protection, ONTAP can use up to three spare disks to replace and reconstruct the data from up to three simultaneously failed disks within the RAID group.

Drive and RAID group information for an ONTAP local tier

Some local tier administration tasks require that you know what types of drives compose the local tier, their size, checksum, and status, whether they are shared with other local tiers, and the size and composition of the RAID groups.

Step

1. Show the drives for the local tier, by RAID group:

storage aggregate show-status aggr_name

The drives are displayed for each RAID group in the local tier.

You can see the RAID type of the drive (data, parity, dparity) in the Position column. If the Position column displays shared, then the drive is shared: if it is an HDD, it is a partitioned disk; if it is an SSD, it is part of a storage pool.

```
cluster1::> storage aggregate show-status nodeA fp 1
Owner Node: cluster1-a
Aggregate: nodeA fp 1 (online, mixed raid type, hybrid) (block checksums)
 Plex: /nodeA_fp_1/plex0 (online, normal, active, pool0)
  RAID Group /nodeA fp 1/plex0/rg0 (normal, block checksums, raid dp)
                                    Usable Physical
    Position Disk Pool Type RPM
                                      Size
                                            Size Status
    _____
                    _____ _____ _____ _____ _____
                  0 SAS
    shared 2.0.1
                             10000 472.9GB 547.1GB (normal)
   shared 2.0.3 0 SAS 10000 472.9GB 547.1GB (normal)
    shared 2.0.5
                   0 SAS
                            10000 472.9GB 547.1GB (normal)
                   0 SAS
   shared 2.0.7
                            10000 472.9GB 547.1GB (normal)
                  0 SAS
    shared 2.0.9
                            10000 472.9GB 547.1GB (normal)
    shared 2.0.11
                     0 SAS 10000 472.9GB 547.1GB (normal)
  RAID Group /nodeA flashpool 1/plex0/rg1
  (normal, block checksums, raid4) (Storage Pool: SmallSP)
                                    Usable Physical
    Position Disk Pool Type RPM Size Size Status
    shared 2.0.13 0 SSD
                             - 186.2GB 745.2GB (normal)
    shared 2.0.12 0 SSD - 186.2GB 745.2GB (normal)
8 entries were displayed.
```

Related information

• storage aggregate show-status

Convert from ONTAP RAID-DP to RAID-TEC

If you want the added protection of triple-parity, you can convert from RAID-DP to RAID-TEC. RAID-TEC is recommended if the size of the disks used in your local tier is greater than 4 TiB.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

Before you begin

The local tier that is to be converted must have a minimum of seven disks.

About this task

• Hard disk drive (HDD) local tiers can be converted from RAID-DP to RAID-TEC. This includes HDD tiers in Flash Pool local tiers.

Steps

1. Verify that the local tier is online and has a minimum of six disks:

storage aggregate show-status -aggregate aggregate name

2. Convert the local tier from RAID-DP to RAID-TEC:

storage aggregate modify -aggregate aggregate_name -raidtype raid_tec

3. Verify that the local tier RAID policy is RAID-TEC:

storage aggregate show aggregate_name

Related information

- · storage aggregate modify
- storage aggregate show-status

Convert from ONTAP RAID-TEC to RAID-DP

If you reduce the size of your local tier and no longer need triple parity, you can convert your RAID policy from RAID-TEC to RAID-DP and reduce the number of disks you need for RAID parity.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

Before you begin

The maximum RAID group size for RAID-TEC is larger than the maximum RAID group size for RAID-DP. If the largest RAID-TEC group size is not within the RAID-DP limits, you cannot convert to RAID-DP.

About this task

To understand the implications of converting between RAID types, refer to the parameters for the storage aggregate modify command.

Steps

1. Verify that the local tier is online and has a minimum of six disks:

storage aggregate show-status -aggregate aggregate name

2. Convert the local tier from RAID-TEC to RAID-DP:

storage aggregate modify -aggregate aggregate_name -raidtype raid dp

3. Verify that the local tier RAID policy is RAID-DP:

storage aggregate show aggregate_name

Related information

- storage aggregate modify
- storage aggregate show-status

Considerations for sizing ONTAP RAID groups

Configuring an optimum RAID group size requires a trade-off of factors. You must decide which factors—speed of RAID rebuild, assurance against risk of data loss due to drive failure, optimizing I/O performance, and maximizing data storage space—are most important for the (local tier) aggregate that you are configuring.

When you create larger RAID groups, you maximize the space available for data storage for the same amount of storage used for parity (also known as the "parity tax"). On the other hand, when a disk fails in a larger RAID group, reconstruction time is increased, impacting performance for a longer period of time. In addition, having more disks in a RAID group increases the probability of a multiple disk failure within the same RAID group.

HDD or array LUN RAID groups

You should follow these guidelines when sizing your RAID groups composed of HDDs or array LUNs:

• All RAID groups in an local tier (aggregate) should have the same number of disks.

While you can have up to 50% less or more than the number of disks in different raid groups on one local tier, this might lead to performance bottlenecks in some cases, so it is best avoided.

• The recommended range of RAID group disk numbers is between 12 and 20.

The reliability of performance disks can support a RAID group size of up to 28, if needed.

• If you can satisfy the first two guidelines with multiple RAID group disk numbers, you should choose the larger number of disks.

SSD RAID groups in Flash Pool local tiers (aggregates)

The SSD RAID group size can be different from the RAID group size for the HDD RAID groups in a Flash Pool local tier (aggregate). Usually, you should ensure that you have only one SSD RAID group for a Flash Pool local tier, to minimize the number of SSDs required for parity.

SSD RAID groups in SSD local tiers (aggregates)

You should follow these guidelines when sizing your RAID groups composed of SSDs:

• All RAID groups in a local tier (aggregate) should have a similar number of drives.

The RAID groups do not have to be exactly the same size, but you should avoid having any RAID group that is less than one half the size of other RAID groups in the same local tier when possible.

• For RAID-DP, the recommended range of RAID group size is between 20 and 28.

Customize the size of your ONTAP RAID groups

You can customize the size of your RAID groups to ensure that your RAID group sizes are appropriate for the amount of storage you plan to include for a local tier.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

About this task

For standard local tiers, you change the size of RAID groups for each local tier separately. For Flash Pool local tiers, you can change the RAID group size for the SSD RAID groups and the HDD RAID groups independently.

The following list outlines some facts about changing the RAID group size:

- By default, if the number of disks or array LUNs in the most recently created RAID group is less than the new RAID group size, disks or array LUNs will be added to the most recently created RAID group until it reaches the new size.
- All other existing RAID groups in that local tier remain the same size, unless you explicitly add disks to them.
- You can never cause a RAID group to become larger than the current maximum RAID group size for the local tier.
- You cannot decrease the size of already created RAID groups.
- The new size applies to all RAID groups in that local tier (or, in the case of a Flash Pool local tier, all RAID groups for the affected RAID group type—SSD or HDD).

Steps

1. Use the applicable command:

If you want to	Enter the following command
Change the maximum RAID group size for the SSD RAID groups of a Flash Pool local tier	storage aggregate modify -aggregate aggr_name -cache-raid-group-size size
Change the maximum size of any other RAID groups	storage aggregate modify -aggregate aggr_name -maxraidsize size

Examples

The following command changes the maximum RAID group size of the local tier n1_a4 to 20 disks or array LUNs:

storage aggregate modify -aggregate n1_a4 -maxraidsize 20

The following command changes the maximum RAID group size of the SSD cache RAID groups of the Flash Pool local tier n1_cache_a2 to 24:

storage aggregate modify -aggregate n1_cache_a2 -cache-raid-group-size 24

Related information

storage aggregate modify

Manage Flash Pool local tiers

Flash Pool ONTAP local tier caching policies

Caching policies for the volumes in a Flash Pool local tier let you deploy Flash as a high performance cache for your working data set while using lower-cost HDDs for less frequently accessed data. If you are providing cache to two or more Flash Pool local tiers, you should use Flash Pool SSD partitioning to share SSDs across the local tiers in the Flash Pool.

Caching policies are applied to volumes that reside in Flash Pool local tiers. You should understand how caching policies work before changing them.

In most cases, the default caching policy of auto is the best caching policy to use. The caching policy should be changed only if a different policy provides better performance for your workload. Configuring the wrong caching policy can severely degrade volume performance; the performance degradation could increase gradually over time.

Caching policies combine a read caching policy and a write caching policy. The policy name concatenates the names of the read caching policy and the write caching policy, separated by a hyphen. If there is no hyphen in the policy name, the write caching policy is none, except for the auto policy.

Read caching policies optimize for future read performance by placing a copy of the data in the cache in addition to the stored data on HDDs. For read caching policies that insert data into the cache for write operations, the cache operates as a *write-through* cache.

Data inserted into the cache by using the write caching policy exists only in cache; there is no copy in HDDs. Flash Pool cache is RAID protected. Enabling write caching makes data from write operations available for reads from cache immediately, while deferring writing the data to HDDs until it ages out of the cache.

If you move a volume from a Flash Pool local tier to a single-tier local tier, it loses its caching policy; if you later move it back to a Flash Pool local tier, it is assigned the default caching policy of auto. If you move a volume between two Flash Pool local tier, the caching policy is preserved.

Change a caching policy

You can use the CLI to change the caching policy for a volume that resides on a Flash Pool local tier by using the -caching-policy parameter with the volume create command.

When you create a volume on a Flash Pool local tier, by default, the auto caching policy is assigned to the volume.

Manage Flash Pool caching policies

Determine whether to modify the ONTAP caching policy of Flash Pool local tiers

You can assign cache-retention policies to volumes in Flash Pool local tiers to determine how long the volume data remains in the Flash Pool cache. However, in some cases changing the cache-retention policy might not impact the amount of time the volume's data remains in the cache.

About this task

If your data meets any of the following conditions, changing your cache-retention policy might not have an impact:

- Your workload is sequential.
- Your workload does not reread the random blocks cached in the solid state drives (SSDs).
- The cache size of the volume is too small.

Steps

The following steps check for the conditions that must be met by the data. The task must be done using the CLI in advanced privilege mode.

1. Use the CLI to view the workload volume:

statistics start -object workload volume

2. Determine the workload pattern of the volume:

```
statistics show -object workload_volume -instance volume-workload -counter
sequential reads
```

3. Determine the hit rate of the volume:

statistics show -object wafl_hya_vvol -instance volume -counter read ops replaced pwercent|wc write blks overwritten percent

4. Determine the Cacheable Read and Project Cache Alloc of the volume:

system node run -node node name wafl awa start aggr name

5. Display the AWA summary:

system node run -node node_name wafl awa print aggr_name

6. Compare the volume's hit rate to the Cacheable Read.

If the hit rate of the volume is greater than the Cacheable Read, then your workload does not reread random blocks cached in the SSDs.

7. Compare the volume's current cache size to the Project Cache Alloc.

If the current cache size of the volume is greater than the Project Cache Alloc, then the size of your volume cache is too small.

Related information

- statistics show
- statistics start

Modify caching policies of ONTAP Flash Pool local tiers

You should modify the caching policy of a volume only if a different caching policy is expected to provide better performance. You can modify the caching policy of a volume

on a Flash Pool local tier.

Before you begin

You must determine whether you want to modify your caching policy.

About this task

In most cases, the default caching policy of auto is the best caching policy that you can use. The caching policy should be changed only if a different policy provides better performance for your workload. Configuring the wrong caching policy can severely degrade volume performance; the performance degradation could increase gradually over time. You should use caution when modifying caching policies. If you experience performance issues with a volume for which the caching policy has been changed, you should return the caching policy to auto.

Step

1. Use the CLI to modify the volume's caching policy:

```
volume modify -volume volume name -caching-policy policy name
```

Example

The following example modifies the caching policy of a volume named vol2 to the policy none:

```
volume modify -volume vol2 -caching-policy none
```

Set the cache-retention policy for ONTAP Flash Pool local tiers

You can assign cache-retention policies to volumes in Flash Pool local tiers. Data in volumes with a high cache-retention policy remains in cache longer and data in volumes with a low cache-retention policy is removed sooner. This increases performance of your critical workloads by making high priority information accessible at a faster rate for a longer period of time.

Before you begin

You should know whether your system has any conditions that might prevent the cache-retention policy from having an impact on how long your data remains in cache.

Steps

Use the CLI in advanced privilege mode to perform the following steps:

1. Change the privilege setting to advanced:

```
set -privilege advanced
```

2. Verify the volume's cache-retention policy:

By default the cache retention policy is "normal".

3. Set the cache-retention policy:

```
volume modify -volume volume_name -vserver vserver_name -caching-policy
policy_name
```

- 4. Verify that the volume's cache-retention policy is changed to the option you selected.
- 5. Return the privilege setting to admin:

set -privilege admin

Flash Pool SSD partitioning for ONTAP Flash Pool local tiers using storage pools

If you are providing cache to two or more Flash Pool local tiers, you should use Flash Pool Solid-State Drive (SSD) partitioning. Flash Pool SSD partitioning allows SSDs to be shared by all the local tiers that use the Flash Pool. This spreads the cost of parity over multiple local tiers, increases SSD cache allocation flexibility, and maximizes SSD performance.

For an SSD to be used in a Flash Pool local tier, the SSD must be placed in a storage pool. You cannot use SSDs that have been partitioned for root-data partitioning in a storage pool. After the SSD is placed in the storage pool, the SSD can no longer be managed as a stand-alone disk and cannot be removed from the storage pool unless you destroy the local tiers associated with the Flash Pool and you destroy the storage pool.

SSD storage pools are divided into four equal allocation units. SSDs added to the storage pool are divided into four partitions and one partition is assigned to each of the four allocation units. The SSDs in the storage pool must be owned by the same HA pair. By default, two allocation units are assigned to each node in the HA pair. Allocation units must be owned by the node that owns the local tier it is serving. If more Flash cache is required for local tiers on one of the nodes, the default number of allocation units can be shifted to decrease the number on one node and increase the number on the partner node.

You use spare SSDs to add to an SSD storage pool. If the storage pool provides allocation units to Flash Pool local tiers owned by both nodes in the HA pair, then the spare SSDs can be owned by either node. However, if the storage pool provides allocation units only to Flash Pool local tiers owned by one of the nodes in the HA pair, then the SSD spares must be owned by that same node.

The following illustration is an example of Flash Pool SSD partitioning. The SSD storage pool provides cache to two Flash Pool local tiers:



Storage pool SP1 is composed of five SSDs and a hot spare SSD. Two of the storage pool's allocation units are allocated to Flash Pool FP1, and two are allocated to Flash Pool FP2. FP1 has a cache RAID type of RAID4. Therefore, the allocation units provided to FP1 contain only one partition designated for parity. FP2 has a cache RAID type of RAID-DP. Therefore, the allocation units provided to FP2 include a parity partition and a double-parity partition.

In this example, two allocation units are allocated to each Flash Pool local tier. However, if one Flash Pool local tier required a larger cache, you could allocate three of the allocation units to that Flash Pool local tier, and only one to the other.

Determine ONTAP Flash Pool candidacy and optimal cache size

Before converting an existing local tier to a Flash Pool local tier, you can determine whether the local tier is I/O bound and the best Flash Pool cache size for your workload and budget. You can also check whether the cache of an existing Flash Pool local tier is sized correctly.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

Before you begin

You should know approximately when the local tier you are analyzing experiences its peak load.

Steps

1. Enter advanced mode:

set advanced

If you need to determine whether an existing local tier would be a good candidate for conversion to a Flash Pool local tier, determine how busy the disks in the local tier are during a period of peak load, and how that is affecting latency:

statistics show-periodic -object disk:raid_group -instance raid_group_name
-counter disk busy|user read latency -interval 1 -iterations 60

You can decide whether reducing latency by adding Flash Pool cache makes sense for this local tier.

The following command shows the statistics for the first RAID group of the local tier "aggr1":

statistics show-periodic -object disk:raid_group -instance /aggr1/plex0/rg0
-counter disk busy|user read latency -interval 1 -iterations 60

3. Start Automated Workload Analyzer (AWA):

```
storage automated-working-set-analyzer start -node node_name -aggregate
aggr name
```

AWA begins collecting workload data for the volumes associated with the specified local tier.

4. Exit advanced mode:

set admin

Allow AWA to run until one or more intervals of peak load have occurred. AWA collects workload statistics for the volumes associated with the specified local tier, and analyzes data for up to one rolling week in duration. Running AWA for more than one week will report only on data collected from the most recent week. Cache size estimates are based on the highest loads seen during the data collection period; the load does not need to be high for the entire data collection period.

5. Enter advanced mode:

set advanced

6. Display the workload analysis:

storage automated-working-set-analyzer show -node node name -instance

7. Stop AWA:

storage automated-working-set-analyzer stop node_name

All workload data is flushed and is no longer available for analysis.

8. Exit advanced mode:

set admin

Related information

• statistics show-periodic

Create an ONTAP Flash Pool local tier using physical SSDs

You create a Flash Pool local tier by enabling the feature on an existing local tier composed of HDD RAID groups, and then adding one or more SSD RAID groups to that local tier. This results in two sets of RAID groups for that local tier: SSD RAID groups (the SSD cache) and HDD RAID groups.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

About this task

After you add an SSD cache to an local tier to create a Flash Pool local tier, you cannot remove the SSD cache to convert the local tier back to its original configuration.

By default, the RAID level of the SSD cache is the same as the RAID level of the HDD RAID groups. You can override this default selection by specifying the raidtype option when you add the first SSD RAID groups.

Before you begin

- You must have identified a valid local tier composed of HDDs to convert to a Flash Pool local tier.
- You must have determined write-caching eligibility of the volumes associated with the local tier, and completed any required steps to resolve eligibility issues.
- You must have determined the SSDs you will be adding, and these SSDs must be owned by the node on which you are creating the Flash Pool local tier.
- You must have determined the checksum types of both the SSDs you are adding and the HDDs already in the local tier.
- You must have determined the number of SSDs you are adding and the optimal RAID group size for the SSD RAID groups.

Using fewer RAID groups in the SSD cache reduces the number of parity disks required, but larger RAID groups require RAID-DP.

- You must have determined the RAID level you want to use for the SSD cache.
- You must have determined the maximum cache size for your system and determined that adding SSD cache to your local tier will not cause you to exceed it.
- You must have familiarized yourself with the configuration requirements for Flash Pool local tiers.

Steps

You can create a Flash Pool local tier using System Manager or the ONTAP CLI.

System Manager

Beginning with ONTAP 9.12.1, you can use System Manager to create a Flash Pool local tier using physical SSDs.

Steps

- 1. Select Storage > Tiers then select an existing local HDD storage tier.
- 2. Select then Add Flash Pool Cache.
- 3. Select Use dedicated SSDs as cache.
- 4. Select a disk type and the number of disks.
- 5. Choose a RAID type.
- 6. Select Save.
- 7. Locate the storage tier then select
- 8. Select More Details. Verify that Flash Pool shows as Enabled.

CLI

Steps

1. Mark the local tier as eligible to become a Flash Pool local tier:

storage aggregate modify -aggregate aggr_name -hybrid-enabled true

If this step does not succeed, determine write-caching eligibility for the target local tier.

- 2. Add the SSDs to the local tier by using the storage aggregate add command.
 - You can specify the SSDs by ID or by using the diskcount and disktype parameters.
 - If the HDDs and the SSDs do not have the same checksum type, or if the local tier is a mixedchecksum local tier, then you must use the checksumstyle parameter to specify the checksum type of the disks you are adding to the local tier.
 - You can specify a different RAID type for the SSD cache by using the raidtype parameter.
 - If you want the cache RAID group size to be different from the default for the RAID type you are using, you should change it now, by using the -cache-raid-group-size parameter.

Related information

- storage aggregate add
- storage aggregate modify

Create a Flash Pool local tier using SSD storage pools

Determine whether an ONTAP Flash Pool local tier is using an SSD storage pool

You can configure a Flash Pool local tier by adding one or more allocation units from an SSD storage pool to an existing HDD local tier.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

You manage Flash Pool local tiers differently when they use SSD storage pools to provide their cache than when they use discrete SSDs.

Step

1. Display the local tier's drives by RAID group:

storage aggregate show-status aggr_name

If the local tier is using one or more SSD storage pools, the value for the Position column for the SSD RAID groups is displayed as Shared, and the name of the storage pool is displayed next to the RAID group name.

Related information

• storage aggregate show-status

Add cache to an ONTAP local tier by creating an SSD storage pool

You can provision cache by converting an existing local tier to a Flash Pool local tier by adding solid state drives (SSDs).

You can create solid state drive (SSD) storage pools to provide SSD cache for two to four Flash Pool local tiers. Flash Pool local tiers enable you to deploy flash as high performance cache for your working data set while using lower-cost HDDs for less frequently accessed data.

About this task

• You must supply a disk list when creating or adding disks to a storage pool.

Storage pools do not support a diskcount parameter.

• The SSDs used in the storage pool should be the same size.

System Manager Use System Manager to add an SSD cache (ONTAP 9.12.1 and later)

Beginning with ONTAP 9.12.1, you can use System Manager to add an SSD cache.



Storage pool options are not available on AFF systems.

Steps

- 1. Click Cluster > Disks and then click Show/Hide.
- 2. Select Type and verify that spare SSDs exist on the cluster.
- 3. Click to Storage > Tiers and click Add Storage Pool.
- 4. Select the disk type.
- 5. Enter a disk size.
- 6. Select the number of disks to add to the storage pool.
- 7. Review the estimated cache size.

Use System Manager to add an SSD cache (ONTAP 9.7 only)



Use the CLI procedure if you are using an ONTAP version later than ONTAP 9.7 or earlier than ONTAP 9.12.1.

Steps

- 1. Click (Return to classic version).
- 2. Click Storage > Aggregates & Disks > Aggregates.
- 3. Select the local tier, and then click **Actions > Add Cache**.
- 4. Select the cache source as "storage pools" or "dedicated SSDs".
- 5. Click (Switch to the new experience).
- 6. Click Storage > Tiers to verify the size of the new local tier.

CLI

Use the CLI to create an SSD storage pool

Steps

1. Determine the names of the available spare SSDs:

storage aggregate show-spare-disks -disk-type SSD

The SSDs used in a storage pool can be owned by either node of an HA pair.

2. Create the storage pool:

storage pool create -storage-pool sp_name -disk-list disk1, disk2,...

3. **Optional:** Verify the newly created storage pool:

Results

After the SSDs are placed into the storage pool, they no longer appear as spares on the cluster, even though the storage provided by the storage pool has not yet been allocated to any Flash Pool caches. You cannot add SSDs to a RAID group as discrete drives; their storage can be provisioned only by using the allocation units of the storage pool to which they belong.

Related information

storage aggregate show

Create an ONTAP Flash Pool local tier using SSD storage pool allocation units

You can configure a Flash Pool local tier by adding one or more allocation units from an SSD storage pool to an existing HDD local tier.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

Beginning with ONTAP 9.12.1, you can use the redesigned System Manager to create a Flash Pool local tier using storage pool allocation units.

Before you begin

- You must have identified a valid local tier composed of HDDs to convert to a Flash Pool local tier.
- You must have determined write-caching eligibility of the volumes associated with the local tier, and completed any required steps to resolve eligibility issues.
- You must have created an SSD storage pool to provide the SSD cache to this Flash Pool local tier.

Any allocation unit from the storage pool that you want to use must be owned by the same node that owns the Flash Pool local tier.

· You must have determined how much cache you want to add to the local tier.

You add cache to the local tier by allocation units. You can increase the size of the allocation units later by adding SSDs to the storage pool if there is room.

• You must have determined the RAID type you want to use for the SSD cache.

After you add a cache to the local tier from SSD storage pools, you cannot change the RAID type of the cache RAID groups.

• You must have determined the maximum cache size for your system and determined that adding SSD cache to your local tier will not cause you to exceed it.

You can see the amount of cache that will be added to the total cache size by using the storage pool show command.

• You must have familiarized yourself with the configuration requirements for Flash Pool local tier.

About this task

If you want the RAID type of the cache to be different from that of the HDD RAID groups, you must specify the cache RAID type when you add the SSD capacity. After you add the SSD capacity to the local tier, you can no longer change the RAID type of the cache.

After you add an SSD cache to a local tier to create a Flash Pool local tier, you cannot remove the SSD cache to convert the local tier back to its original configuration.

System Manager

Beginning with ONTAP 9.12.1, you can use System Manager to add SSDs to an SSD storage pool.

Steps

- 1. Click Storage > Tiers and select an existing local HDD storage tier.
- 2. Click and select Add Flash Pool Cache.
- 3. Select Use Storage Pools.
- 4. Select a storage pool.
- 5. Select a cache size and RAID configuration.
- 6. Click Save.
- 7. Locate the storage tier again and click
- 8. Select More Details and verify that the Flash Pool shows as Enabled.

CLI

Steps

1. Mark the local tier as eligible to become a Flash Pool local tier:

storage aggregate modify -aggregate aggr_name -hybrid-enabled true

If this step does not succeed, determine write-caching eligibility for the target local tier.

2. Show the available SSD storage pool allocation units:

storage pool show-available-capacity

3. Add the SSD capacity to the local tier:

```
storage aggregate add aggr_name -storage-pool sp_name -allocation-units
number_of_units
```

If you want the RAID type of the cache to be different from that of the HDD RAID groups, you must change it when you enter this command by using the raidtype parameter.

You do not need to specify a new RAID group; ONTAP automatically puts the SSD cache into separate RAID groups from the HDD RAID groups.

You cannot set the RAID group size of the cache; it is determined by the number of SSDs in the storage pool.

The cache is added to the local tier and the local tier is now a Flash Pool local tier. Each allocation unit added to the local tier becomes its own RAID group.

4. Confirm the presence and size of the SSD cache:

storage aggregate show aggregate_name

The size of the cache is listed under Total Hybrid Cache Size.

Related information

- NetApp Technical Report 4070: Flash Pool Design and Implementation Guide
- storage aggregate add
- storage aggregate modify

Determine the impact on ONTAP cache size when SSDs are added to an SSD storage pool

If adding SSDs to a storage pool causes your platform model's cache limit to be exceeded, ONTAP does not allocate the newly added capacity to any Flash Pool local tiers. This can result in some or all of the newly added capacity being unavailable for use.

About this task

When you add SSDs to an SSD storage pool that has allocation units already allocated to Flash Pool local tiers, you increase the cache size of each of those local tiers and the total cache on the system. If none of the storage pool's allocation units have been allocated, adding SSDs to that storage pool does not affect the SSD cache size until one or more allocation units are allocated to a cache.

Steps

1. Determine the usable size of the SSDs you are adding to the storage pool:

storage disk show disk_name -fields usable-size

2. Determine how many allocation units remain unallocated for the storage pool:

storage pool show-available-capacity sp_name

All unallocated allocation units in the storage pool are displayed.

- 3. Calculate the amount of cache that will be added by applying the following formula:
 - (4 number of unallocated allocation units) × 25% × usable size × number of SSDs

Add SSDs to an ONTAP SSD storage pool

When you add solid state drives (SSDs) to an SSD storage pool, you increase the storage pool's physical and usable sizes and allocation unit size. The larger allocation unit size also affects allocation units that have already been allocated to local tiers.

Before you begin

You must have determined that this operation will not cause you to exceed the cache limit for your HA pair. ONTAP does not prevent you from exceeding the cache limit when you add SSDs to an SSD storage pool, and doing so can render the newly added storage capacity unavailable for use.

About this task

When you add SSDs to an existing SSD storage pool, the SSDs must be owned by one node or the other of the same HA pair that already owned the existing SSDs in the storage pool. You can add SSDs that are owned by either node of the HA pair.

The SSD you add to the storage pool must be the same size as disk currently used in the storage pool.

System Manager

Beginning with ONTAP 9.12.1, you can use System Manager to add SSDs to an SSD storage pool.

Steps

- 1. Click **Storage > Tiers** and locate the **Storage Pools** section.
- 2. Locate the storage pool, click , and select Add Disks.
- 3. Choose the disk type and select the number of disks.
- 4. Review the estimate cache size.

CLI

Steps

1. **Optional:** View the current allocation unit size and available storage for the storage pool:

storage pool show -instance *sp_name*

2. Find available SSDs:

storage disk show -container-type spare -type SSD

3. Add the SSDs to the storage pool:

storage pool add -storage-pool sp name -disk-list disk1, disk2 ...

The system displays which Flash Pool local tiers will have their size increased by this operation and by how much, and prompts you to confirm the operation.

ONTAP commands for managing SSD storage pools

ONTAP provides the storage pool command for managing SSD storage pools.

If you want to	Use this command
Display how much storage a storage pool is providing to which local tier	storage pool show-aggregate
Display how much cache would be added to the overall cache capacity for both RAID types (allocation unit data size)	storage pool show -instance
Display the disks in a storage pool	storage pool show-disks
Display the unallocated allocation units for a storage pool	storage pool show-available-capacity
Change the ownership of one or more allocation units of a storage pool from one HA partner to the other	storage pool reassign

Related information

ONTAP command reference

FabricPool tier management

Learn about data tiering with ONTAP FabricPool

You can use FabricPool to automatically tier data depending on how frequently the data is accessed.

FabricPool is a hybrid storage solution that on AFF systems uses an all flash (all SSD) aggregate, and on FAS systems uses either an all flash (all SSD) or HDD aggregate as the performance tier and an object store as the cloud tier. Using a FabricPool helps you reduce storage cost without compromising performance, efficiency, or protection.

The cloud tier can be located on NetApp StorageGRID or ONTAP S3 (beginning with ONTAP 9.8), or one of the following service providers:

- Alibaba cloud
- Amazon S3
- Amazon Commercial Cloud Services
- Google Cloud
- IBM cloud
- Microsoft Azure Blob Storage



Beginning with ONTAP 9.7, additional object store providers that support generic S3 APIs can be used by selecting the S3_Compatible object store provider.

Related information

• NetApp cloud tiering documentation

Requirements for using ONTAP FabricPool

To help ensure that you optimize your FabricPool configurations, you should familiarize yourself with a few considerations and requirements about using FabricPool.

General considerations and requirements

ONTAP 9.4

- You must be running ONTAP 9.4 or later releases for the following FabricPool functionality:
 - The auto tiering policy
 - $\circ\,$ Specifying the tiering minimum cooling period
 - Inactive data reporting (IDR)
 - $\circ\,$ Using Microsoft Azure Blob Storage for the cloud as the cloud tier for FabricPool
 - Using FabricPool with ONTAP Select

ONTAP 9.5

- You must be running ONTAP 9.5 or later releases for the following FabricPool functionality:
 - Specifying the tiering fullness threshold
 - · Using IBM Cloud Object Storage as the cloud tier for FabricPool
 - NetApp Volume Encryption (NVE) of the cloud tier, enabled by default.

ONTAP 9.6

- You must be running ONTAP 9.6 or later releases for the following FabricPool functionality:
 - The all tiering policy
 - Inactive data reporting enabled manually on HDD aggregates
 - Inactive data reporting enabled automatically for SSD aggregates when you upgrade to ONTAP 9.6 and at time aggregate is created, except on low end systems with less than 4 CPU, less than 6 GB of RAM, or when WAFL-buffer-cache size is less than 3 GB.

ONTAP monitors system load, and if the load remains high for 4 continuous minutes, IDR is disabled, and is not automatically enabled. You can reenable IDR manually; however, manually enabled IDR is not automatically disabled.

- · Using Alibaba Cloud Object Storage as the cloud tier for FabricPool
- Using Google Cloud Platform as the cloud tier for FabricPool
- Volume move without cloud tier data copy

ONTAP 9.7

- You must be running ONTAP 9.7 or later releases for the following FabricPool functionality:
 - Non transparent HTTP and HTTPS proxy to provide access to only whitelisted access points, and to provide auditing and reporting capabilities.
 - · FabricPool mirroring to tier cold data to two object stores simultaneously
 - · FabricPool mirrors on MetroCluster configurations
 - NDMP dump and restore enabled by default on FabricPool attached aggregates.



If the backup application uses a protocol other than NDMP, such as NFS or SMB, all data being backed up in the performance tier becomes hot and can affect tiering of that data to the cloud tier. Non-NDMP reads can cause data migration from the cloud tier back to the performance tier.

NDMP Backup and Restore Support for FabricPool

ONTAP 9.8

- You must be running ONTAP 9.8 or later for the following FabricPool functionality:
 - Cloud retrieval
 - FabricPool with SnapLock Enterprise. FabricPool with SnapLock Enterprise requires a Feature Product Variance Request (FPVR). To create an FPVR, contact your sales team.
 - Minimum cooling period maximum of 183 days

- · Object tagging using user-created custom tags
- HDD FabricPool aggregates

HDD FabricPools are supported with SAS, FSAS, BSAS and MSATA disks only on systems with 6 or more CPU cores.

Check Hardware Universe for the latest supported models.

ONTAP 9.10.1

- You must be running ONTAP 9.10.1 or later for the following FabricPool functionality:
 - PUT throttling
 - Temperature-sensitive storage efficiency (TSSE).

ONTAP 9.12.1

- You must be running ONTAP 9.12.1 or later for the following FabricPool functionality:
 - SVM Migrate
 - Support for FabricPool, FlexGroup, and SVM-DR working in conjunction. (Prior to 9.12.1 any two of these features worked together, but not all three in conjunction.)

ONTAP 9.14.1

- You must be running ONTAP 9.14.1 or later for the following FabricPool functionality:
 - Cloud Write
 - Aggressive Readahead

Local tiers (aggregates)

FabricPool supports the following aggregate types:

- On AFF systems, you can only use SSD aggregates for FabricPool.
- On FAS systems, you can use either SSD or HDD aggregates for FabricPool.
- On Cloud Volumes ONTAP and ONTAP Select, you can use either SSD or HDD aggregates for FabricPool. Using SSD aggregates is recommended.



Flash Pool aggregates, which contain both SSDs and HDDs, are not supported.

Cloud tiers

FabricPool supports using the following object stores as the cloud tier:

- Alibaba Cloud Object Storage Service (Standard, Infrequent Access)
- Amazon S3 (Standard, Standard-IA, One Zone-IA, Intelligent-Tiering, Glacier Instant Retrieval)
- Amazon Commercial Cloud Services (C2S)
- Google Cloud Storage (Multi-Regional, Regional, Nearline, Coldline, Archive)
- IBM Cloud Object Storage (Standard, Vault, Cold Vault, Flex)

- Microsoft Azure Blob Storage (Hot and Cool)
- NetApp ONTAP S3 (ONTAP 9.8 and later)
- NetApp StorageGRID (StorageGRID 10.3 and later)



Glacier Flexible Retrieval and Glacier Deep Archive are not supported.

- The object store "bucket" (container) you plan to use must have already been set up, must have at least 10 GB of storage space, and must not be renamed.
- You cannot detach a cloud tier from a local tier after it is attached; however, you can use FabricPool mirror to attach a local tier to a different cloud tier.

Intercluster LIFs

Cluster high-availability (HA) pairs that use FabricPool require two intercluster LIFs to communicate with the cloud tier. NetApp recommends creating an intercluster LIF on additional HA pairs to seamlessly attach cloud tiers to local tiers on those nodes as well.

Disabling or deleting an intercluster LIF interrupts communication to the cloud tier.



Because concurrent SnapMirror and SnapVault replication operations share the network link to the cloud tier, initialization and RTO are dependent on the available bandwidth and latency to the cloud tier. Performance degradation might occur if connectivity resources become saturated. Proactive configuration of multiple LIFs can significantly decrease this type of network saturation.

If you are using more than one intercluster LIF on a node with different routing, NetApp recommends placing them in different IPspaces. During configuration, FabricPool can select from multiple IPspaces, but it is unable to select specific intercluster LIFs within an IPspace.

Network Time Protocol (NTP)

Network Time Protocol (NTP) configuration is required to ensure the time is synchronized between clusters. Learn about how to configure NTP.

ONTAP storage efficiencies

Storage efficiencies such as compression, deduplication, and compaction are preserved when moving data to the cloud tier, reducing required object storage capacity and transport costs.



Beginning with ONTAP 9.15.1, FabricPool supports Intel QuickAssist Technology (QAT4) which provides more aggressive, and more performant, storage efficiency savings.

Aggregate inline deduplication is supported on the local tier, but associated storage efficiencies are not carried over to objects stored on the cloud tier.

When using the All volume tiering policy, storage efficiencies associated with background deduplication processes might be reduced as data is likely to be tiered before the additional storage efficiencies can be applied.

BlueXP tiering license

FabricPool requires a capacity-based license when attaching third-party object storage providers (such as Amazon S3) as cloud tiers for AFF and FAS systems. A BlueXP Tiering license is not required when using StorageGRID or ONTAP S3 as the cloud tier or when tiering with Cloud Volumes ONTAP, Amazon FSx for NetApp ONTAP, or Azure NetApp files.

BlueXP licenses (including add-on or extensions to preexisting FabricPool licenses) are activated in the BlueXP digital wallet.

StorageGRID consistency controls

StorageGRID's consistency controls affects how the metadata that StorageGRID uses to track objects is distributed between nodes and the availability of objects for client requests. NetApp recommends using the default, read-after-new-write, consistency control for buckets used as FabricPool targets.



Do not use the available consistency control for buckets used as FabricPool targets.

Additional considerations for tiering data accessed by SAN protocols

When tiering data that is accessed by SAN protocols, NetApp recommends using private clouds, like ONTAP S3 or StorageGRID, due to connectivity considerations.



You should be aware that when using FabricPool in a SAN environment with a Windows host, if the object storage becomes unavailable for an extended period of time when tiering data to the cloud, files on the NetApp LUN on the Windows host might become inaccessible or disappear. See the Knowledge Base article During FabricPool S3 object store unavailable Windows SAN host reported filesystem corruption.

Quality of Service

• If you use throughput floors (QoS Min), the tiering policy on the volumes must be set to none before the aggregate can be attached to FabricPool.

Other tiering policies prevent the aggregate from being attached to FabricPool. A QoS policy will not enforce throughput floors when FabricPool is enabled.

Functionality or features not supported by FabricPool

- Object stores with WORM enabled and object versioning enabled.
- Information lifecycle management (ILM) policies that are applied to object store buckets

FabricPool supports StorageGRID's Information Lifecycle Management policies only for data replication and erasure coding to protect cloud tier data from failure. However, FabricPool does *not* support advanced ILM rules such as filtering based on user metadata or tags. ILM typically includes various movement and deletion policies. These policies can be disruptive to the data in the cloud tier of FabricPool. Using FabricPool with ILM policies that are configured on object stores can result in data loss.

- 7-Mode data transition using the ONTAP CLI commands or the 7-Mode Transition Tool
- RAID SyncMirror, except in a MetroCluster configuration
- SnapLock volumes when using ONTAP 9.7 and earlier releases

Tamperproof snapshots

Tamperproof snapshots provide immutable protections that cannot be deleted. Because FabricPool requires the ability to delete data, FabricPool and snapshot locks cannot be enabled on the same volume.

- Tape backup using SMTape for FabricPool-enabled aggregates
- The Auto Balance functionality
- Volumes using a space guarantee other than none

With the exception of root SVM volumes and CIFS audit staging volumes, FabricPool does not support attaching a cloud tier to an aggregate that contains volumes using a space guarantee other than none. For example, a volume using a space guarantee of volume (-space-guarantee volume) is not supported.

- Clusters with DP_Optimized license
- Flash Pool aggregates

Tier data efficiently with ONTAP FabricPool policies

FabricPool tiering policies enable you to move data efficiently across tiers as data becomes hot or cold. Understanding the tiering policies helps you select the right policy that suits your storage management needs.

Types of FabricPool tiering policies

FabricPool tiering policies determine when or whether the user data blocks of a volume in FabricPool are moved to the cloud tier, based on the volume "temperature" of hot (active) or cold (inactive). The volume "temperature" increases when it is accessed frequently and decreases when it is not. Some tiering policies have an associated tiering minimum cooling period, which sets the time that user data in a volume of FabricPool must remain inactive for the data to be considered "cold" and moved to the cloud tier.

After a block has been identified as cold, it is marked as eligible to be tiered. A daily background tiering scan looks for cold blocks. When enough 4KB blocks from the same volume have been collected, they are concatenated into a 4MB object and moved to the cloud tier based on the volume tiering policy.



Data in volumes using the all tiering policy is immediately marked as cold and begins tiering to the cloud tier as soon as possible. It does not need to wait for the daily tiering scan to run.

You can use the volume object-store tiering show command to view the tiering status of a FabricPool volume.

Learn more about volume object-store tiering show in the ONTAP command reference.

The FabricPool tiering policy is specified at the volume level. Four options are available:

• The snapshot-only tiering policy (the default) moves user data blocks of the volume snapshots that are not associated with the active file system to the cloud tier.

The tiering minimum cooling period is 2 days. You can modify the default setting for the tiering minimum cooling period with the -tiering-minimum-cooling-days parameter in the advanced privilege level of the volume create and volume modify commands. Valid values are 2 to 183 days using ONTAP 9.8 and later. If you are using a version of ONTAP earlier than 9.8, valid values are 2 to 63 days.

• The auto tiering policy, supported only on ONTAP 9.4 and later releases, moves cold user data blocks in both the snapshots and the active file system to the cloud tier.

The default tiering minimum cooling period is 31 days and applies to the entire volume, for both the active file system and the snapshots.

You can modify the default setting for the tiering minimum cooling period with the -tiering-minimum -cooling-days parameter in the advanced privilege level of the volume create and volume modify commands. Valid values are 2 to 183 days.

• The all tiering policy, supported only with ONTAP 9.6 and later, moves all user data blocks in both the active file system and snapshots to the cloud tier. It replaces the backup tiering policy.

The all volume tiering policy should not be used on read/write volumes that have normal client traffic.

The tiering minimum cooling period does not apply because the data moves to the cloud tier as soon as the tiering scan runs, and you cannot modify the setting.

• The none tiering policy keeps a volume's data in the performance tier and does not move cold to the cloud tier.

Setting the tiering policy to none prevents new tiering. Volume data that has previously been moved to the cloud tier remains in the cloud tier until it becomes hot and is automatically moved back to the local tier.

The tiering minimum cooling period does not apply because the data never moves to the cloud tier, and you cannot modify the setting.

When cold blocks in a volume with a tiering policy set to none are read, they are made hot and written to the local tier.

The volume show command output shows the tiering policy of a volume. A volume that has never been used with FabricPool shows the none tiering policy in the output.



When in an SVM DR relationship, source and destination volumes do not need to use FabricPool aggregates, but they must use the same tiering policy.

What happens when you modify the tiering policy of a volume in FabricPool

You can modify the tiering policy of a volume by performing a volume modify operation. You must understand how changing the tiering policy might affect how long it takes for data to become cold and be moved to the cloud tier.

- Changing the tiering policy from snapshot-only or none to auto causes ONTAP to send user data blocks in the active file system that are already cold to the cloud tier, even if those user data blocks were not previously eligible for the cloud tier.
- Changing the tiering policy to all from another policy causes ONTAP to move all user blocks in the active file system and in the snapshots to the cloud as soon as possible. Prior to ONTAP 9.8, blocks needed to wait until the next tiering scan ran.

Moving blocks back to the performance tier is not allowed.

• Changing the tiering policy from auto to snapshot-only or none does not cause active file system

blocks that are already moved to the cloud tier to be moved back to the performance tier.

Volume reads are needed for the data to be moved back to the performance tier.

• Any time you change the tiering policy on a volume, the tiering minimum cooling period is reset to the default value for the policy.

What happens to the tiering policy when you move a volume

• Unless you explicitly specify a different tiering policy, a volume retains its original tiering policy when it is moved in and out of a FabricPool-enabled aggregate.

However, the tiering policy takes effect only when the volume is in a FabricPool-enabled aggregate.

• The existing value of the -tiering-minimum-cooling-days parameter for a volume moves with the volume unless you specify a different tiering policy for the destination.

If you specify a different tiering policy, then the volume uses the default tiering minimum cooling period for that policy. This is the case whether the destination is FabricPool or not.

- You can move a volume across aggregates and at the same time modify the tiering policy.
- You should pay special attention when a volume move operation involves the auto tiering policy.

Assuming that both the source and the destination are FabricPool-enabled aggregates, the following table summarizes the outcome of a volume move operation that involves policy changes related to auto:

When you move a volume that has a tiering policy of	And you change the tiering policy with the move to	Then after the volume move
all	auto	All data is moved to the performance tier.
snapshot-only, none, or auto	auto	Data blocks are moved to the same tier of the destination as they previously were on the source.
auto or all	snapshot-only	All data is moved to the performance tier.
auto	all	All user data is moved to the cloud tier.
snapshot-only,auto or all	none	All data is kept at the performance tier.

What happens to the tiering policy when you clone a volume

• Beginning with ONTAP 9.8, a clone volume always inherits both the tiering policy and the cloud retrieval policy from the parent volume.

In releases earlier than ONTAP 9.8, a clone inherits the tiering policy from the parent except when the parent has the all tiering policy.

- If the parent volume has the never cloud retrieval policy, its clone volume must have either the never cloud retrieval policy or the all tiering policy, and a corresponding cloud retrieval policy default.
- The parent volume cloud retrieval policy cannot be changed to never unless all its clone volumes have a cloud retrieval policy never.

When you clone volumes, keep the following best practices in mind:

- The -tiering-policy option and tiering-minimum-cooling-days option of the clone only controls the tiering behavior of blocks unique to the clone. Therefore, we recommend using tiering settings on the parent FlexVol that are either move the same amount of data or move less data than any of the clones
- The cloud retrieval policy on the parent FlexVol should either move the same amount of data or should move more data than the retrieval policy of any of the clones

How tiering policies work with cloud migration

FabricPool cloud data retrieval is controlled by tiering policies that determine data retrieval from the cloud tier to performance tier based on the read pattern. Read patterns can be either sequential or random.

The following table lists the tiering policies and the cloud data retrieval rules for each policy.

Tiering policy	Retrieval behavior
none	Sequential and random reads
snapshot-only	Sequential and random reads
auto	Random reads
all	No data retrieval

Beginning with ONTAP 9.8, the cloud migration control cloud-retrieval-policy option overrides the default cloud migration or retrieval behavior controlled by the tiering policy.

The following table lists the supported cloud retrieval policies and their retrieval behavior.

Cloud retrieval policy	Retrieval behavior
default	Tiering policy decides what data should be pulled back, so there is no change to cloud data retrieval with "default," cloud-retrieval-policy. This policy is the default value for any volume regardless of the hosted aggregate type.
on-read	All client-driven data read is pulled from cloud tier to performance tier.

never	No client-driven data is pulled from cloud tier to performance tier
promote	 For tiering policy "none," all cloud data is pulled from the cloud tier to the performance tier For tiering policy "snapshot-only," AFS data is pulled.

Learn more about the commands described in this procedure in the ONTAP command reference.

Learn about ONTAP FabricPool configuration and management tasks

You can use the FabricPool workflow diagram to help you plan the configuration and management tasks.



Configure FabricPool

Prepare for FabricPool configuration

Get started with ONTAP FabricPool

Configuring FabricPool helps you manage which storage tier (the local performance tier or the cloud tier) data should be stored based on whether the data is frequently accessed.

The preparation required for FabricPool configuration depends on the object store you use as the cloud tier.

The FabricPool license you might have used in the past is changing and is being retained only for configurations that aren't supported within BlueXP. Beginning August 21, 2021, BlueXP tiering BYOL licensing was introduced for tiering configurations that are supported within BlueXP using the BlueXP tiering service.

Learn more about BlueXP tiering BYOL licensing.

Configurations that are supported by BlueXP must use the Digital Wallet page in BlueXP to license tiering for ONTAP clusters. This requires you to set up a BlueXP account and set up tiering for the particular object storage provider you plan to use. BlueXP currently supports tiering to the following object storage: Amazon S3, Azure Blob storage, Google Cloud Storage, S3-compatible object storage, and StorageGRID.

Learn more about the BlueXP tiering service.

You can download and activate a FabricPool license using System Manager if you have one of the configurations that is not supported within BlueXP:

- ONTAP installations in Dark Sites
- ONTAP clusters that are tiering data to IBM Cloud Object Storage or Alibaba Cloud Object Storage

The FabricPool license is a cluster-wide license. It includes an entitled usage limit that you purchase for object storage that is associated with FabricPool in the cluster. The usage across the cluster must not exceed the capacity of the entitled usage limit. If you need to increase the usage limit of the license, you should contact your sales representative.

FabricPool licenses are available in perpetual or term-based, 1- or 3- year, formats.

A term-based FabricPool license with 10 TB of free capacity is available for first time FabricPool orders for existing clusters configurations not supported within BlueXP. Free capacity is not available with perpetual licenses.

A license is not required if you use NetApp StorageGRID or ONTAP S3 for the cloud tier. Cloud Volumes ONTAP does not require a FabricPool license, regardless of the provider you are using.

This task is supported only by uploading the license file to the cluster using System Manager.

Steps

- 1. Download the NetApp License File (NLF) for the FabricPool license from the NetApp Support Site.
- 2. Perform the following actions using System Manager to upload the FabricPool license to the cluster:
 - a. In the Cluster > Settings pane, on the Licenses card, click ->.
 - b. On the License page, click + Add.
 - c. In the **Add License** dialog box, click **Browse** to select the NLF you downloaded, and then click **Add** to upload the file to the cluster.

Related information

ONTAP FabricPool (FP) Licensing Overview

NetApp Software License Search

NetApp TechComm TV: FabricPool playlist

Using CA certificates creates a trusted relationship between client applications and StorageGRID.

Unless you plan to disable certificate checking for StorageGRID, you must install a StorageGRID CA certificate on the cluster so that ONTAP can authenticate with StorageGRID as the object store for FabricPool.

Although StorageGRID can generate self-signed certificates, using signed certificates from a third-party certificate authority is the recommended best practice.

About this task

Although installation and use of certificate authority (CA) certificates are recommended best practices, beginning with ONTAP 9.4, installation of CA certificates is not required for StorageGRID.

Steps

- 1. Contact your StorageGRID administrator to obtain the StorageGRID system's CA certificate.
- 2. Use the security certificate install command with the -type server-ca parameter to install the StorageGRID CA certificate on the cluster.

The fully qualified domain name (FQDN) you enter must match the custom common name on the StorageGRID CA certificate.

Update an expired certificate

To update an expired certificate, the best practice is to use a trusted CA to generate the new server certificate. In addition, you should ensure that the certificate is updated on the StorageGRID server and on the ONTAP cluster at the same time to keep any downtime to a minimum.

Related information

- StorageGRID Resources
- security certificate install

Install a CA certificate on a cluster for ONTAP S3

Using CA certificates creates a trusted relationship between client applications and the ONTAP S3 object store server. A CA certificate should be installed on ONTAP before using it as an object store that is accessible to remote clients.

Unless you plan to disable certificate checking for ONTAP S3, you must install a ONTAP S3 CA certificate on the cluster so that ONTAP can authenticate with ONTAP S3 as the object store for FabricPool.

Although ONTAP can generate self-signed certificates, using signed certificates from a third-party certificate authority is the recommended best practice.

Steps

- 1. Obtain the ONTAP S3 system's CA certificate.
- 2. Use the security certificate install command with the -type server-ca parameter to install the ONTAP S3 CA certificate on the cluster.

The fully qualified domain name (FQDN) you enter must match the custom common name on the ONTAP

Update an expired certificate

To update an expired certificate, the best practice is to use a trusted CA to generate the new server certificate. In addition, you should ensure that the certificate is updated on the ONTAP S3 server and on the ONTAP cluster at the same time to keep any downtime to a minimum.

You can use System Manager to renew an expired certificate on an ONTAP cluster.

Steps

- 1. Navigate to **Cluster > Settings**.
- 2. Scroll to the Security section, locate the Certificates pane, and click ->.
- 3. In the Trusted certificate authorities tab, locate the name of the certificate you want to renew.
- 4. Next to the certificate name click and select **Renew**.
- 5. In the **Renew trusted certificate authority** window, copy and paste or import the certificate information into the **Certificate details** area.
- 6. Click Renew.

Related information

- S3 configuration
- security certificate install

Set up an object store as the cloud tier for FabricPool

Set up an object store as the cloud tier for FabricPool overview

Setting up FabricPool involves specifying the configuration information of the object store (StorageGRID, ONTAP S3, Alibaba Cloud Object Storage, Amazon S3, Google Cloud Storage, IBM Cloud Object Storage, or Microsoft Azure Blob Storage for the cloud) that you plan to use as the cloud tier for FabricPool.

Set up StorageGRID as the ONTAP FabricPool cloud tier

You can set up StorageGRID as the cloud tier for FabricPool. When tiering data that is accessed by SAN protocols, NetApp recommends using private clouds, like StorageGRID, due to connectivity considerations.

Considerations for using StorageGRID with FabricPool

- You need to install a CA certificate for StorageGRID, unless you explicitly disable certificate checking.
- Do not enable StorageGRID object versioning on the object store bucket.
- A FabricPool license is not required.
- If a StorageGRID node is deployed in a virtual machine with storage assigned from a NetApp AFF system, confirm that the volume does not have a FabricPool tiering policy enabled.

Disabling FabricPool tiering for volumes used with StorageGRID nodes simplifies troubleshooting and storage operations.


Never use FabricPool to tier any data related to StorageGRID back to StorageGRID itself. Tiering StorageGRID data back to StorageGRID increases troubleshooting and operational complexity.

About this task

Load balancing is enabled for StorageGRID in ONTAP 9.8 and later. When the server's hostname resolves to more than one IP address, ONTAP establishes client connections with all the IP addresses returned (up to a maximum of 16 IP addresses). The IP addresses are picked up in a round-robin method when connections are established.

Steps

You can set up StorageGRID as the cloud tier for FabricPool with ONTAP System Manager or the ONTAP CLI.

- 1. Click Storage > Tiers > Add Cloud Tier and select StorageGRID as the object store provider.
- 2. Complete the requested information.
- 3. If you want to create a cloud mirror, click Add as FabricPool Mirror.

A FabricPool mirror provides a method for you to seamlessly replace a data store, and it helps to ensure that your data is available in the event of disaster.

CLI

- 1. Specify the StorageGRID configuration information by using the storage aggregate objectstore config create command with the -provider-type SGWS parameter.
 - The storage aggregate object-store config create command fails if ONTAP cannot access StorageGRID with the provided information.
 - You use the -access-key parameter to specify the access key for authorizing requests to the StorageGRID object store.
 - You use the -secret-password parameter to specify the password (secret access key) for authenticating requests to the StorageGRID object store.
 - If the StorageGRID password is changed, you should update the corresponding password stored in ONTAP immediately.

Doing so enables ONTAP to access the data in StorageGRID without interruption.

 Setting the -is-certificate-validation-enabled parameter to false disables certificate checking for StorageGRID. Using signed certificates (-is-certificate-validation -enabled true) from a third-party certificate authority is a recommended best practice.

```
cluster1::> storage aggregate object-store config create
-object-store-name mySGWS -provider-type SGWS -server mySGWSserver
-container-name mySGWScontainer -access-key mySGWSkey
-secret-password mySGWSpass
```

2. Display and verify the StorageGRID configuration information by using the storage aggregate object-store config show command.

The storage aggregate object-store config modify command enables you to modify the StorageGRID configuration information for FabricPool.

Related information

- storage aggregate object-store config create
- storage aggregate object-store config modify
- storage aggregate object-store config show

Set up ONTAP S3 as the FabricPool cloud tier

If you are running ONTAP 9.8 or later, you can set up ONTAP S3 as the cloud tier for FabricPool.

Before you begin

• You must have the ONTAP S3 server name and the IP address of its associated LIFs on the remote cluster.



The server name is used as the fully qualified domain name (FQDN) by client applications. Outside of ONTAP, confirm DNS records point to the SVM data LIFs being used.

• There must be intracluster LIFs on the local cluster.

When configured for local cluster tiering, a local tier (also known as a storage aggregate in the ONTAP CLI) is attached to a local bucket. FabricPool uses cluster LIFs for intracluster traffic.



Performance degradation might occur if cluster LIF resources become saturated. To avoid this, NetApp recommends using four-node or greater clusters when tiering to a local bucket along with an HA pair for the local tier and an HA pair for the local bucket. Tiering to local buckets on a single HA pair is not recommended.

• To enable remote FabricPool capacity (cloud) tiering using ONTAP S3, you must configure intercluster LIFs on the FabricPool client and configure data LIFs on the object store server.

About this task

Load balancing is enabled for ONTAP S3 servers in ONTAP 9.8 and later. When the server's hostname resolves to more than one IP address, ONTAP establishes client connections with all the IP addresses returned (up to a maximum of 16 IP addresses). The IP addresses are picked up in a round-robin method when connections are established.

Steps

You can set up ONTAP S3 as the cloud tier for FabricPool with ONTAP System Manager or the ONTAP CLI.

- 1. Click Storage > Tiers > Add Cloud Tier and select ONTAP S3 as the object store provider.
- 2. Complete the requested information.
- 3. If you want to create a cloud mirror, click Add as FabricPool Mirror.

A FabricPool mirror provides a method for you to seamlessly replace a data store, and it helps to ensure that your data is available in the event of disaster.

CLI

1. Add entries for the S3 server and LIFs to your DNS server.

Option	Description
If you use an external DNS server	Give the S3 server name and IP addresses to the DNS server administrator.
If you use your local system's DNS hosts table	Enter the following command: dns host create -vserver <svm_name> -address ip_address -hostname <s3_server_name></s3_server_name></svm_name>

- 2. Specify the ONTAP S3 configuration information by using the storage aggregate objectstore config create command with the -provider-type ONTAP S3 parameter.
 - The storage aggregate object-store config create command fails if the local ONTAP system cannot access the ONTAP S3 server with the information provided.
 - You use the -access-key parameter to specify the access key for authorizing requests to the ONTAP S3 server.
 - You use the -secret-password parameter to specify the password (secret access key) for authenticating requests to the ONTAP S3 server.
 - If the ONTAP S3 server password is changed, you should immediately update the corresponding password stored in the local ONTAP system.

Doing so enables access to the data in the ONTAP S3 object store without interruption.

• Setting the -is-certificate-validation-enabled parameter to false disables certificate checking for ONTAP S3. Using signed certificates (-is-certificate-validation-enabled true) from a third-party certificate authority is a recommended best practice.

```
cluster1::> storage aggregate object-store config create
-object-store-name myS3 -provider-type ONTAP_S3 -server myS3server
-container-name myS3container -access-key myS3key
-secret-password myS3pass
```

3. Display and verify the ONTAP_S3 configuration information by using the storage aggregate object-store config show command.

The storage aggregate object-store config modify command enables you to modify the ONTAP S3 configuration information for FabricPool.

Related information

- Create LIF for SMB
- Create LIF for NFS
- storage aggregate object-store config create
- storage aggregate object-store config modify
- storage aggregate object-store config show

Set up Alibaba Cloud Object Storage as the ONTAP FabricPool cloud tier

If you are running ONTAP 9.6 or later, you can set up Alibaba Cloud Object Storage as the cloud tier for FabricPool.

Considerations for using Alibaba Cloud Object Storage with FabricPool

- A BlueXP tiering license is required when tiering to Alibaba Cloud Object Storage. For more information, see Install a FabricPool license on an ONTAP cluster.
- On AFF and FAS systems and ONTAP Select, FabricPool supports the following Alibaba Object Storage Service classes:
 - · Alibaba Object Storage Service Standard
 - Alibaba Object Storage Service Infrequent Access

Alibaba Cloud: Introduction to storage classes

Contact your NetApp sales representative for information about storage classes not listed.

Steps

- 1. Specify the Alibaba Cloud Object Storage configuration information by using the storage aggregate object-store config create command with the -provider-type AliCloud parameter.
 - The storage aggregate object-store config create command fails if ONTAP cannot access Alibaba Cloud Object Storage with the provided information.
 - You use the -access-key parameter to specify the access key for authorizing requests to the Alibaba Cloud Object Storage object store.
 - If the Alibaba Cloud Object Storage password is changed, you should update the corresponding password stored in ONTAP immediately.

Doing so enables ONTAP to access the data in Alibaba Cloud Object Storage without interruption.

```
storage aggregate object-store config create my_ali_oss_store_1
-provider-type AliCloud -server oss-us-east-1.aliyuncs.com
-container-name my-ali-oss-bucket -access-key DXJRXHPXHYXA9X31X3JX
```

2. Display and verify the Alibaba Cloud Object Storage configuration information by using the storage aggregate object-store config show command.

The storage aggregate object-store config modify command enables you to modify the Alibaba Cloud Object Storage configuration information for FabricPool.

Related information

- storage aggregate object-store config create
- storage aggregate object-store config modify
- storage aggregate object-store config show

Set up Amazon S3 as the ONTAP FabricPool cloud tier

You can set up Amazon S3 as the cloud tier for FabricPool. If you are running ONTAP 9.5 or later, you can set up Amazon Commercial Cloud Services (C2S) for FabricPool.

Considerations for using Amazon S3 with FabricPool

- A BlueXP tiering license is required when tiering to Amazon S3.
- It is recommended that the LIF that ONTAP uses to connect with the Amazon S3 object server be on a 10 Gbps port.
- On AFF and FAS systems and ONTAP Select, FabricPool supports the following Amazon S3 storage classes:
 - Amazon S3 Standard
 - · Amazon S3 Standard Infrequent Access (Standard IA)
 - · Amazon S3 One Zone Infrequent Access (One Zone IA)
 - Amazon S3 Intelligent-Tiering
 - Amazon Commercial Cloud Services
 - Beginning with ONTAP 9.11.1, Amazon S3 Glacier Instant Retrieval (FabricPool does not support Glacier Flexible Retrieval or Glacier Deep Archive)

Amazon Web Services Documentation: Amazon S3 Storage Classes

Contact your sales representative for information about storage classes not listed.

• On Cloud Volumes ONTAP, FabricPool supports tiering from General Purpose SSD (gp2) and Throughput Optimized HDD (st1) volumes of Amazon Elastic Block Store (EBS).

Steps

- 1. Specify the Amazon S3 configuration information by using the storage aggregate object-store config create command with the -provider-type AWS_S3 parameter.
 - You use the -auth-type CAP parameter to obtain credentials for C2S access.

When you use the *-auth-type CAP* parameter, you must use the *-cap-url* parameter to specify the full URL to request temporary credentials for C2S access.

- The storage aggregate object-store config create command fails if ONTAP cannot access Amazon S3 with the provided information.
- You use the -access-key parameter to specify the access key for authorizing requests to the Amazon S3 object store.
- You use the -secret-password parameter to specify the password (secret access key) for authenticating requests to the Amazon S3 object store.
- If the Amazon S3 password is changed, you should update the corresponding password stored in ONTAP immediately.

Doing so enables ONTAP to access the data in Amazon S3 without interruption.

```
cluster1::> storage aggregate object-store config create
-object-store-name my_aws_store -provider-type AWS_S3
-server s3.amazonaws.com -container-name my-aws-bucket
-access-key DXJRXHPXHYXA9X31X3JX
```

```
cluster1::> storage aggregate object-store config create -object
-store-name my_c2s_store -provider-type AWS_S3 -auth-type CAP -cap
-url
https://123.45.67.89/api/v1/credentials?agency=XYZ&mission=TESTACCT&r
ole=S3FULLACCESS -server my-c2s-s3server-fqdn -container my-c2s-s3-
bucket
```

 Display and verify the Amazon S3 configuration information by using the storage aggregate objectstore config show command.

The storage aggregate object-store config modify command enables you to modify the Amazon S3 configuration information for FabricPool.

Related information

- · storage aggregate object-store config create
- · storage aggregate object-store config modify
- storage aggregate object-store config show

Set up Google Cloud Storage as the ONTAP FabricPool cloud tier

If you are running ONTAP 9.6 or later, you can set up Google Cloud Storage as the cloud tier for FabricPool.

Additional considerations for using Google Cloud Storage with FabricPool

• A BlueXP tiering license is required when tiering to Google Cloud Storage.

- It is recommended that the LIF that ONTAP uses to connect with the Google Cloud Storage object server be on a 10 Gbps port.
- On AFF and FAS systems and ONTAP Select, FabricPool supports the following Google Cloud Object storage classes:
 - Google Cloud Multi-Regional
 - Google Cloud Regional
 - · Google Cloud Nearline
 - Google Cloud Coldline

Google Cloud: Storage Classes

Steps

- 1. Specify the Google Cloud Storage configuration information by using the storage aggregate objectstore config create command with the -provider-type GoogleCloud parameter.
 - The storage aggregate object-store config create command fails if ONTAP cannot access Google Cloud Storage with the provided information.
 - You use the -access-key parameter to specify the access key for authorizing requests to the Google Cloud Storage object store.
 - If the Google Cloud Storage password is changed, you should update the corresponding password stored in ONTAP immediately.

Doing so enables ONTAP to access the data in Google Cloud Storage without interruption.

```
storage aggregate object-store config create my_gcp_store_1 -provider
-type GoogleCloud -container-name my-gcp-bucket1 -access-key
GOOGAUZZUV2USCFGHGQ51118
```

2. Display and verify the Google Cloud Storage configuration information by using the storage aggregate object-store config show command.

The storage aggregate object-store config modify command enables you to modify the Google Cloud Storage configuration information for FabricPool.

Related information

- storage aggregate object-store config create
- storage aggregate object-store config modify
- storage aggregate object-store config show

Set up IBM Cloud Object Storage as the ONTAP FabricPool cloud tier

If you are running ONTAP 9.5 or later, you can set up IBM Cloud Object Storage as the cloud tier for FabricPool.

Considerations for using IBM Cloud Object Storage with FabricPool

- A BlueXP tiering license is required when tiering to IBM Cloud Object Storage.
- It is recommended that the LIF that ONTAP uses to connect with the IBM Cloud object server be on a 10 Gbps port.

Steps

- 1. Specify the IBM Cloud Object Storage configuration information by using the storage aggregate object-store config create command with the -provider-type IBM COS parameter.
 - The storage aggregate object-store config create command fails if ONTAP cannot access IBM Cloud Object Storage with the provided information.
 - You use the -access-key parameter to specify the access key for authorizing requests to the IBM Cloud Object Storage object store.
 - You use the -secret-password parameter to specify the password (secret access key) for authenticating requests to the IBM Cloud Object Storage object store.
 - If the IBM Cloud Object Storage password is changed, you should update the corresponding password stored in ONTAP immediately.

Doing so enables ONTAP to access the data in IBM Cloud Object Storage without interruption.

```
storage aggregate object-store config create
-object-store-name MyIBM -provider-type IBM_COS
-server s3.us-east.objectstorage.softlayer.net
-container-name my-ibm-cos-bucket -access-key DXJRXHPXHYXA9X31X3JX
```

2. Display and verify the IBM Cloud Object Storage configuration information by using the storage aggregate object-store config show command.

The storage aggregate object-store config modify command enables you to modify the IBM Cloud Object Storage configuration information for FabricPool.

Related information

- storage aggregate object-store config create
- storage aggregate object-store config modify
- storage aggregate object-store config show

Set up Azure Blob Storage as the ONTAP FabricPool cloud tier

If you are running ONTAP 9.4 or later, you can set up Azure Blob Storage as the cloud tier for FabricPool.

Considerations for using Microsoft Azure Blob Storage with FabricPool

- A BlueXP tiering license is required when tiering to Azure Blob Storage.
- A FabricPool license is not required if you are using Azure Blob Storage with Cloud Volumes ONTAP.
- It is recommended that the LIF that ONTAP uses to connect with the Azure Blob Storage object server be on a 10 Gbps port.
- FabricPool currently does not support Azure Stack, which is on-premises Azure services.

• At the account level in Microsoft Azure Blob Storage, FabricPool supports only hot and cool storage tiers.

FabricPool does not support blob-level tiering. It also does not support tiering to Azure's archive storage tier.

About this task

FabricPool currently does not support Azure Stack, which is on-premises Azure services.

Steps

- 1. Specify the Azure Blob Storage configuration information by using the storage aggregate objectstore config create command with the -provider-type Azure Cloud parameter.
 - The storage aggregate object-store config create command fails if ONTAP cannot access Azure Blob Storage with the provided information.
 - You use the -azure-account parameter to specify the Azure Blob Storage account.
 - You use the -azure-private-key parameter to specify the access key for authenticating requests to Azure Blob Storage.
 - If the Azure Blob Storage password is changed, you should update the corresponding password stored in ONTAP immediately.

Doing so enables ONTAP to access the data in Azure Blob Storage without interruption.

```
cluster1::> storage aggregate object-store config create
-object-store-name MyAzure -provider-type Azure_Cloud
-server blob.core.windows.net -container-name myAzureContainer
-azure-account myAzureAcct -azure-private-key myAzureKey
```

 Display and verify the Azure Blob Storage configuration information by using the storage aggregate object-store config show command.

The storage aggregate object-store config modify command enables you to modify the Azure Blob Storage configuration information for FabricPool.

Related information

- storage aggregate object-store config create
- storage aggregate object-store config modify
- storage aggregate object-store config show

Set up object stores for ONTAP FabricPool in a MetroCluster configuration

If you are running ONTAP 9.7 or later, you can set up a mirrored FabricPool on a MetroCluster configuration to tier cold data to object stores in two different fault zones.

About this task

• FabricPool in MetroCluster requires that the underlying mirrored aggregate and the associated object store configuration must be owned by the same MetroCluster configuration.

- You cannot attach an aggregate to an object store that is created in the remote MetroCluster site.
- You must create object store configurations on the MetroCluster configuration that owns the aggregate.

Before you begin

- The MetroCluster configuration is set up and properly configured.
- Two objects stores are set up on the appropriate MetroCluster sites.
- · Containers are configured on each of the object stores.
- IP spaces are created or identified on the two MetroCluster configurations and their names match.

Step

1. Specify the object store configuration information on each MetroCluster site by using the storage object-store config create command.

In this example, FabricPool is required on only one cluster in the MetroCluster configuration. Two object store configurations are created for that cluster, one for each object store bucket.

```
storage aggregate
   object-store config create -object-store-name mccl-ostore-config-s1
-provider-type SGWS -server
   <SGWS-server-1> -container-name <SGWS-bucket-1> -access-key <key>
-secret-password <password> -encrypt
        <true|false> -provider <provider-type> -is-ssl-enabled <true|false>
ipspace
        <IPSpace>
```

```
storage aggregate object-store config create -object-store-name mccl-
ostore-config-s2
    -provider-type SGWS -server <SGWS-server-2> -container-name <SGWS-
bucket-2> -access-key <key> -secret-password <password> -encrypt
<true|false> -provider <provider-type>
    -is-ssl-enabled <true|false> ipspace <IPSpace>
```

This example sets up FabricPool on the second cluster in the MetroCluster configuration.

```
storage aggregate
   object-store config create -object-store-name mcc2-ostore-config-s1
-provider-type SGWS -server
   <SGWS-server-1> -container-name <SGWS-bucket-3> -access-key <key>
-secret-password <password> -encrypt
        <true|false> -provider <provider-type> -is-ssl-enabled <true|false>
ipspace
        <IPSpace>
```

```
storage aggregate
   object-store config create -object-store-name mcc2-ostore-config-s2
-provider-type SGWS -server
   <SGWS-server-2> -container-name <SGWS-bucket-4> -access-key <key>
-secret-password <password> -encrypt
        <true|false> -provider <provider-type> -is-ssl-enabled <true|false>
ipspace
        <IPSpace>
```

Related information

• storage object-store config create

Test the ONTAP cloud tier latency and throughput performance

Before you attach an object store to a local tier, you can test the object store's latency and throughput performance by using object store profiler.

Object store profiler results are a measurement of connectivity between ONTAP and the cloud tier object store using 4MB PUTs and random-read byte-ranged GETs ranging from 4MB to 256KB. (Only internal ONTAP features, such as SnapMirror, can make use of GETs larger than 32KB.)

Because they do not account for competing workloads or unique client application behavior, object store profiler results are not a perfect indicator of tiering performance.

Before you begin

1

- You must add the cloud tier to ONTAP before you can use it with the object store profiler.
- You must be at the ONTAP CLI advanced privilege mode.

Steps

1. Start the object store profiler:

```
storage aggregate object-store profiler start -object-store-name <name> -node
<name>
```

2. View the results:

storage aggregate object-store profiler show

Related information

- storage aggregate object-store profiler show
- storage aggregate object-store profiler start

Associate the ONTAP cloud tier with a local tier

After setting up an object store as the cloud tier, you specify the local tier to use by attaching it to FabricPool. In ONTAP 9.5 and later, you can also attach local tiers that

contain qualified FlexGroup volume constituents.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

About this task

Attaching a cloud tier to a local tier is a permanent action. A cloud tier cannot be unattached from a local tier after being attached. However, you can use FabricPool mirror to attach a local tier to a different cloud tier.

Before you begin

When you use the ONTAP CLI to set up an local tier for FabricPool, the local tier must already exist.



When you use System Manager to set up a local tier for FabricPool, you can create the local tier and set it up to use for FabricPool at the same time.

Steps

You can attach a local tier to a FabricPool object store with ONTAP System Manager or the ONTAP CLI.

- 1. Navigate to **Storage > Tiers**, select a cloud tier, then click
- 2. Select Attach local tiers.
- 3. Under Add as Primary verify that the volumes are eligible to attach.
- 4. If necessary, select Convert volumes to thin provisioned.
- 5. Click Save.

CLI

To attach an object store to an aggregate with the CLI:

1. **Optional**: To see how much data in a volume is inactive, follow the steps in Determining how much data in a volume is inactive by using inactive data reporting.

Seeing how much data in a volume is inactive can help you decide which aggregate to use for FabricPool.

2. Attach the object store to an aggregate by using the storage aggregate object-store attach command.

If the aggregate has never been used with FabricPool and it contains existing volumes, then the volumes are assigned the default snapshot-only tiering policy.

```
cluster1::> storage aggregate object-store attach -aggregate myaggr
-object-store-name Amazon01B1
```

You can use the allow-flexgroup true option to attach aggregates that contain FlexGroup volume constituents.

3. Display the object store information and verify that the attached object store is available by using the storage aggregate object-store show command.

Related information

- storage aggregate object-store attach
- storage aggregate object-store show

Tier data to a local ONTAP S3 bucket

Beginning with ONTAP 9.8, you can tier data to local object storage using ONTAP S3.

Tiering data to a local bucket provides a simple alternative to moving data to a different local tier. This procedure uses either an existing bucket on the local cluster, or you can let ONTAP automatically create a new storage VM and a new bucket.

Keep in mind that once you attach the primary local bucket it cannot be unattached.

Before you begin

- An S3 license is required for this workflow, which creates a new S3 server and new bucket, or uses existing ones. This license is included in ONTAP One. A FabricPool license is not required for this workflow.
- Enable ONTAP S3 access for local FabricPool tiering.

Steps

- 1. Tier data to a local bucket: click **Storage > Tiers**, in the **SSD** pane, select a local tier, click ;, and select **Tier to local bucket**.
- 2. In the Primary tier section, choose either Existing or New.
- 3. Click Save.

Manage FabricPool

Analyze inactive ONTAP data with inactive data reporting

Seeing how much data in a volume is inactive enables you to make good use of storage tiers. Information in inactive data reporting helps you decide which aggregate to use for FabricPool, whether to move a volume in to or out of FabricPool, or whether to modify the tiering policy of a volume.

Before you begin

You must be running ONTAP 9.4 or later to use the inactive data reporting functionality.

About this task

• Inactive data reporting is not supported on some aggregates.

You cannot enable inactive data reporting when FabricPool cannot be enabled, including the following instances:

- Root aggregates
- MetroCluster aggregates running ONTAP versions earlier than 9.7
- Flash Pool (hybrid aggregates, or SnapLock aggregates)
- Inactive data reporting is enabled by default on aggregates where any volumes have adaptive compression enabled.
- Inactive data reporting is enabled by default on all SSD aggregates in ONTAP 9.6.
- Inactive data reporting is enabled by default on FabricPool aggregate in ONTAP 9.4 and ONTAP 9.5.
- You can enable inactive data reporting on non-FabricPool aggregates using the ONTAP CLI, including HDD aggregates, beginning with ONTAP 9.6.

Procedure

You can determine how much data is inactive with ONTAP System Manager or the ONTAP CLI.

1. Choose one of the following options:

- When you have existing HDD aggregates, navigate to Storage > Tiers and click i for the aggregate on which you want to enable inactive data reporting.
- When no cloud tiers are configured, navigate to **Dashboard** and click the **Enable inactive data reporting** link under **Capacity**.

CLI

To enable inactive data reporting with the CLI:

1. If the aggregate for which you want to see inactive data reporting is not used in FabricPool, enable inactive data reporting for the aggregate by using the storage aggregate modify command with the -is-inactive-data-reporting-enabled true parameter.

```
cluster1::> storage aggregate modify -aggregate aggr1 -is-inactive
-data-reporting-enabled true
```

You need to explicitly enable the inactive data reporting functionality on an aggregate that is not used for FabricPool.

You cannot and do not need to enable inactive data reporting on a FabricPool-enabled aggregate because the aggregate already comes with inactive data reporting. The <code>-is-inactive-data</code> -reporting-enabled parameter does not work on FabricPool-enabled aggregates.

The -fields is-inactive-data-reporting-enabled parameter of the storage aggregate show command shows whether inactive data reporting is enabled on an aggregate.

2. To display how much data is inactive on a volume, use the volume show command with the -fields performance-tier-inactive-user-data, performance-tier-inactive-user-data-percent parameter.

cluster1::> volume show -fields performance-tier-inactive-user- data,performance-tier-inactive-user-data-percent			
vserver	volume	performance-tier-inactive-user-data	performance-tier-
inactive-user-data-percent			
vsim1	volO	0B	0%
vs1	vs1rv1	0B	0%
vs1	vv1	10.34MB	0%
vs1	vv2	10.38MB	0%
4 entrie	es were	displayed.	

• The performance-tier-inactive-user-data field displays how much user data stored in the aggregate is inactive.

- The performance-tier-inactive-user-data-percent field displays what percent of the data is inactive across the active file system and snapshots.
- For an aggregate that is not used for FabricPool, inactive data reporting uses the tiering policy to decide how much data to report as cold.
 - For the none tiering policy, 31 days is used.
 - For the snapshot-only and auto, inactive data reporting uses tiering-minimum-cooling-days.
 - For the ALL policy, inactive data reporting assumes the data will tier within a day.

Until the period is reached, the output shows "-" for the amount of inactive data instead of a value.

- On a volume that is part of FabricPool, what ONTAP reports as inactive depends on the tiering policy that is set on a volume.
 - For the none tiering policy, ONTAP reports the amount of the entire volume that is inactive for at least 31 days. You cannot use the -tiering-minimum-cooling-days parameter with the none tiering policy.
 - For the ALL, snapshot-only, and auto tiering policies, inactive data reporting is not supported.

Related information

• storage aggregate modify

Manage volumes for FabricPool

Create a volume on a FabricPool-enabled ONTAP local tier

You can add volumes to FabricPool by creating new volumes directly in the FabricPoolenabled local tier or by moving existing volumes from another local tier to the FabricPoolenabled local tier.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

When you create a volume for FabricPool, you have the option to specify a tiering policy. If no tiering policy is specified, the created volume uses the default <code>snapshot-only</code> tiering policy. For a volume with the <code>snapshot-only</code> or <code>auto</code> tiering policy, you can also specify the tiering minimum cooling period.

Before you begin

- Setting a volume to use the auto tiering policy or specifying the tiering minimum cooling period requires ONTAP 9.4 or later.
- Using FlexGroup volumes requires ONTAP 9.5 or later.
- Setting a volume to use the all tiering policy requires ONTAP 9.6 or later.
- Setting a volume to use the -cloud-retrieval-policy parameter requires ONTAP 9.8 or later.

Steps

- 1. Create a new volume for FabricPool by using the volume create command.
 - The -tiering-policy optional parameter enables you to specify the tiering policy for the volume.

You can specify one of the following tiering policies:

- snapshot-only (default)
- auto
- all
- backup (deprecated)
- none

Types of FabricPool tiering policies

• The -cloud-retrieval-policy optional parameter enables cluster administrators with the advanced privilege level to override the default cloud migration or retrieval behavior controlled by the tiering policy.

You can specify one of the following cloud retrieval policies:

• default

The tiering policy determines what data is pulled back, so there is no change to cloud data retrieval with default cloud-retrieval-policy. This means the behavior is the same as in pre-ONTAP 9.8 releases:

- If the tiering policy is none or snapshot-only, then "default" means that any client-driven data read is pulled from the cloud tier to performance tier.
- If the tiering policy is auto, then any client-driven random read is pulled but not sequential reads.
- If the tiering policy is all then no client-driven data is pulled from the cloud tier.
- on-read

All client-driven data reads are pulled from the cloud tier to performance tier.

never

No client-driven data is pulled from the cloud tier to performance tier

- promote
 - For tiering policy none, all cloud data is pulled from the cloud tier to the performance tier
 - For tiering policy snapshot-only, all active filesystem data is pulled from the cloud tier to the performance tier.

• The -tiering-minimum-cooling-days optional parameter in the advanced privilege level enables you to specify the tiering minimum cooling period for a volume that uses the snapshot-only or auto tiering policy.

Beginning with ONTAP 9.8, you can specify a value between 2 and 183 for the tiering minimum cooling days. If you are using a version of ONTAP earlier than 9.8, you can specify a value between 2 and 63

for the tiering minimum cooling days.

Example of creating a volume for FabricPool

The following example creates a volume called "myvol1" in the "myFabricPool" FabricPool-enabled local tier. The tiering policy is set to auto and the tiering minimum cooling period is set to 45 days:

```
cluster1::*> volume create -vserver myVS -aggregate myFabricPool
-volume myvol1 -tiering-policy auto -tiering-minimum-cooling-days 45
```

Related information

FlexGroup volumes management

Move a volume to a FabricPool-enabled ONTAP local tier

A volume move is the way that ONTAP moves a volume nondisruptively from one local tier (source) to another (destination). Volume moves can be performed for a variety of reasons, although the most common reasons are hardware lifecycle management, cluster expansion, and load balancing.

It is important to understand how volume move works with FabricPool because the changes that take place at both the local tier, the attached cloud tier, and the volume (volume tiering policies) can have a major impact on functionality.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

Destination local tier

If a volume move's destination local tier does not have an attached cloud tier, data on the source volume that is stored on the cloud tier is written to the local tier on the destination local tier.

Beginning with ONTAP 9.8, when a volume has inactive data reporting enabled, FabricPool will use the volume's heat map to immediately queue cold data to begin tiering as soon as it is written to the destination local tier.

Prior to ONTAP 9.8, moving a volume to another local tier resets the inactivity period of blocks on the local tier. For example, a volume using the Auto volume tiering policy with data on the local tier that has been inactive for 20 days, but had not yet tiered, will have the temperature of the data reset to 0 days after a volume move.

Optimized volume moves

Beginning with ONTAP 9.6, if a volume move's destination local tier uses the same bucket as the source local tier, data on the source volume that is stored in the bucket does not move back to the local tier. Tiered data stays at rest and only hot data needs to be moved from one local tier to another. This optimized volume move results in significant network efficiencies.

For example, a 300TB optimized volume move means that even though 300TB of cold data moves from one local tier to another, it will not trigger 300TB of reads and 300TB of writes to the object store.

Unoptimized volume moves generate additional network and compute traffic (reads/GETs and writes/PUTs), increasing demands on the ONTAP cluster and object store, potentially raising costs when tiering to public object stores.

Some configurations are incompatible with optimized volume moves:

- · Changing tiering policy during volume move
- · Source and destination local tiers using different encryption keys
- FlexClone volumes
- FlexClone parent volumes
- MetroCluster (supports optimized volume moves in ONTAP 9.8 and later)
- Unsynchronized FabricPool Mirror buckets

If a volume move's destination local tier has an attached cloud tier, data on the source volume that is stored on the cloud tier is first written to the local tier on the destination local tier. It is then written to the cloud tier on the destination local tier if this approach is appropriate for the volume's tiering policy.

Writing data to the local tier first improves the performance of the volume move and reduces cutover time. If a volume tiering policy is not specified when performing a volume move, the destination volume uses the tiering policy of the source volume.

If a different tiering policy is specified when performing the volume move, the destination volume is created with the specified tiering policy and the volume move is not optimized.

Volume metadata

Regardless of whether a volume move is optimized, ONTAP stores a significant amount of metadata about the location, storage efficiency, permissions, usage patterns, etc., of all data, both local and tiered. Metadata always stays on the local tier and is not tiered. When a volume is moved from one local tier to another, this information needs to be moved to the destination local tier as well.

Duration

Volume moves still take time to complete and the expectation should be that an optimized volume move will take approximately the same amount of time as moving an equal amount of non-tiered data.

It is important to understand that "throughput" reported by the volume move show command does not represent throughput in terms of data being moved from the cloud tier, but volume data being updated locally.



When in an SVM DR relationship, source and destination volumes must use the same tiering policy.

Steps

1. Use the volume move start command to move a volume from a source local tier to a destination local tier.

Example of moving a volume

The following example moves a volume named $\tt myvol2$ of <code>vsl SVM</code> to <code>dest_FabricPool</code>, a FabricPool enabled local tier.

```
cluster1::> volume move start -vserver vs1 -volume myvol2
-destination-aggregate dest_FabricPool
```

Enable ONTAP volumes in FabricPool to write directly to the cloud

Beginning with ONTAP 9.14.1, you can enable and disable writing directly to the cloud on a new or existing volume in a FabricPool to allow NFS clients to write data directly to the cloud without waiting for tiering scans. SMB clients still write to the performance tier in a cloud write enabled volume. Cloud-write mode is disabled by default.

Having the ability to write directly to the cloud is helpful for cases like migrations, for example, where large amounts of data are transferred to a cluster than the cluster can support on the local tier. Without cloud write mode, during a migration, smaller amounts of data are transferred, then tiered, then transferred and tiered again, until the migration is complete. Using cloud write mode, this type of management is no longer required because the data is never transferred to the local tier.

Before you begin

- You should be a cluster or SVM administrator.
- You must be at the advanced privilege level.
- The volume must be a read-write type volume.
- The volume must have the ALL tiering policy.

Enable writing directly to the cloud during volume creation

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. Create a volume and enable cloud write mode:

```
volume create -vserver <svm name> -volume <volume name> -is-cloud-write
-enabled <true|false> -aggregate <local tier name>
```

The following example creates a volume named vol1 with cloud write enabled on the FabricPool local tier (aggr1):

```
volume create -vserver vs1 -volume vol1 -is-cloud-write-enabled true
-aggregate aggr1
```

Enable writing directly to the cloud on an existing volume

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. Modify a volume to enable cloud write mode:

```
volume modify -vserver <svm name> -volume <volume name> -is-cloud-write
-enabled true
```

The following example modifies the volume named vol1 to enable cloud write:

volume modify -vserver vs1 -volume vol1 -is-cloud-write-enabled true

Disable writing directly to the cloud on a volume

Steps

1. Set the privilege level to advanced:

set -privilege advanced

2. Disable cloud write mode on a volume:

```
volume modify -vserver <svm name> -volume <volume name> -is-cloud-write
-enabled false
```

The following example disables cloud write mode on the volume named vol1:

volume modify -vserver vs1 -volume vol1 -is-cloud-write-enabled false

Enable ONTAP volumes in FabricPool to perform aggressive read-aheads

Beginning with ONTAP 9.14.1, you can enable and disable aggressive read-ahead mode on volumes in FabricPools. In ONTAP 9.13.1, aggressive read-ahead mode was introduced only on cloud platforms. Beginning with ONTAP 9.14.1, aggressive readahead mode is available on all platforms that FabricPool supports, including on-premises platforms. The feature is disabled by default.

When aggressive read-ahead is *disabled*, FabricPool only reads the file blocks that a client application needs; it does not need to read the entire file. This can result in reduced network traffic, especially for large GB-sized and TB-sized files. *Enabling* aggressive read-ahead on a volume turns this functionality off, and FabricPool preemptively reads the entire file sequentially from the object store, increasing GET throughput and reducing

the latency of client reads on the file. By default, when tiered data is read sequentially it stays cold and is not written to the local tier.

Aggressive read-ahead trades network efficiency for increased performance of tiered data.

About this task

The aggressive-readahead-mode command has two options:

- none: read-ahead is disabled.
- file prefetch: the system reads the entire file into memory ahead of the client application.

Before you begin

- You should be a cluster or SVM administrator.
- You must be at the advanced privilege level.

Enable aggressive read-ahead mode during volume creation

Steps

1. Set the privilege level to advanced:

set -privilege advanced

2. Create a volume and enable aggressive read-ahead mode:

```
volume create -volume <volume name> -aggressive-readahead-mode
<none|file prefetch>
```

The following example creates a volume named vol1 with aggressive read-ahead enabled with the file_prefetch option:

volume create -volume vol1 -aggressive-readahead-mode file prefetch

Disable aggressive read-ahead mode

Steps

1. Set the privilege level to advanced:

set -privilege advanced

2. Disable aggressive read-ahead mode:

```
volume modify -volume <volume name> -aggressive-readahead-mode none
```

The following example modifies a volume named vol1 to disable aggressive read-ahead mode:

volume modify -volume vol1 -aggressive-readahead-mode none

View aggressive read-ahead mode on a volume

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. View the aggressive read-ahead mode:

```
volume show -fields aggressive-readahead-mode
```

Manage ONTAP FabricPool volumes with user-created custom tags

Beginning with ONTAP 9.8, FabricPool supports object tagging using user-created custom tags to enable you to classify and sort objects for easier management. If you are a user with the admin privilege level, you can create new object tags, and modify, delete, and view existing tags.

Assign a new tag during volume creation

You can create a new object tag when you want to assign one or more tags to new objects that are tiered from a new volume you create. You can use tags to help you classify and sort tiering objects for easier data management. Beginning with ONTAP 9.8, you can use System Manager to create object tags.

About this task

You can set tags only on FabricPool volumes attached to StorageGRID. These tags are retained during a volume move.

- A maximum of four tags per volume is allowed.
- In the CLI, each object tag must be a key-value pair separated by an equal sign.
- In the CLI, multiple tags must be separated by a comma.
- Each tag value can contain a maximum of 127 characters.
- Each tag key must start with either an alphabetic character or an underscore.

Keys must contain only alphanumeric characters and underscores, and the maximum number of characters allowed is 127.

You can assign object tags with ONTAP System Manager or the ONTAP CLI.

- 1. Navigate to **Storage > Tiers**.
- 2. Locate a storage tier with volumes you want to tag.
- 3. Click the Volumes tab.
- 4. Locate the volume you want to tag and in the **Object Tags** column select **Click to enter tags**.
- 5. Enter a key and value.
- 6. Click Apply.

CLI

1. Use the volume create command with the -tiering-object-tags option to create a new volume with the specified tags. You can specify multiple tags in comma-separated pairs:

```
volume create [ -vserver <vserver name> ] -volume <volume_name>
-tiering-object-tags <key1=value1>
[,<key2=value2>,<key3=value3>,<key4=value4> ]
```

The following example creates a volume named fp_volume1 with three object tags.

vol create -volume fp_volume1 -vserver vs0 -tiering-object-tags
project=fabricpool,type=abc,content=data

Modify an existing tag

You can change the name of a tag, replace tags on existing objects in the object store, or add a different tag to new objects that you plan to add later.

- 1. Navigate to **Storage > Tiers**.
- 2. Locate a storage tier with volumes containing tags you want to modify.
- 3. Click the Volumes tab.
- 4. Locate the volume with tags you want to modify, and in the **Object Tags** column click the tag name.
- 5. Modify the tag.
- 6. Click Apply.

CLI

1. Use the volume modify command with the -tiering-object-tags option to modify an existing tag.

```
volume modify [ -vserver <vserver name> ] -volume <volume_name>
-tiering-object-tags <key1=value1> [ ,<key2=value2>,
<key3=value3>,<key4=value4> ]
```

The following example changes the name of the existing tag type=abc to type=xyz.

vol modify -volume fp_volume1 -vserver vs0 -tiering-object-tags
project=fabricpool,type=xyz,content=data

Delete a tag

You can delete object tags when you no longer want them set on a volume or on objects in the object store.

- 1. Navigate to **Storage > Tiers**.
- 2. Locate a storage tier with volumes containing tags you want to delete.
- 3. Click the Volumes tab.
- 4. Locate the volume with tags you want to delete, and in the **Object Tags** column click the tag name.
- 5. To delete the tag, click the trash can icon.
- 6. Click Apply.

CLI

1. Use the volume modify command with the -tiering-object-tags option followed by an empty value ("") to delete an existing tag.

The following example deletes the existing tags on fp_volume1.

vol modify -volume fp volume1 -vserver vs0 -tiering-object-tags ""

View existing tags on a volume

You can view the existing tags on a volume to see what tags are available before appending new tags to the list.

Steps

1. Use the volume show command with the tiering-object-tags option to view existing tags on a volume.

```
volume show [ -vserver <vserver name> ] -volume <volume_name> -fields
tiering-object-tags
```

Check object tagging status on FabricPool volumes

You can check if tagging is complete on one or more FabricPool volumes.

Steps

1. Use the vol show command with the -fields needs-object-retagging option to see if tagging is in progress, if it has completed, or if tagging is not set.

```
vol show -fields needs-object-retagging [ -instance | -volume <volume
name>]
```

One of the following values is displayed:

- true: the object tagging scanner has not yet to run or needs to run again for this volume
- ° false: the object tagging scanner has completed tagging for this volume
- <->: the object tagging scanner is not applicable for this volume. This happens for volumes that are not residing on FabricPools.

Monitor space utilization of a FabricPool-enabled ONTAP local tier

You need to know how much data is stored in the performance and cloud tiers for FabricPool. That information helps you determine whether you need to change the tiering policy of a volume, increase the FabricPool licensed usage limit, or increase the storage space of the cloud tier.



Prior to ONTAP 9.7, System Manager uses the term *aggregate* to describe a *local tier*. Regardless of your ONTAP version, the ONTAP CLI uses the term *aggregate*. To learn more about local tiers, see Disks and local tiers.

Steps

1. Monitor the space utilization for FabricPool-enabled local tiers by using one of the following commands to display the information:

If you want to display	Then use this command:
The used size of the cloud tier in a local tier	storage aggregate show with the -instance parameter
Details of space utilization within an local tiers, including the object store's referenced capacity	storage aggregate show-space with the -instance parameter
Space utilization of the object stores that are attached to the local tiers, including how much license space is being used	storage aggregate object-store show- space
A list of volumes in a local tier and the footprints of their data and metadata	volume show-footprint

In addition to using CLI commands, you can use Active IQ Unified Manager (formerly OnCommand Unified Manager), along with FabricPool Advisor, which is supported on ONTAP 9.4 and later clusters, or System Manager to monitor the space utilization.

The following example shows ways of displaying space utilization and related information for FabricPool:

cluster1::> storage aggregate show-space -instance Aggregate: MyFabricPool . . . Aggregate Display Name: MyFabricPool . . . Total Object Store Logical Referenced Capacity: -Object Store Logical Referenced Capacity Percentage: -. . . Object Store Size: -Object Store Space Saved by Storage Efficiency: -Object Store Space Saved by Storage Efficiency Percentage: -Total Logical Used Size: -Logical Used Percentage: -Logical Unreferenced Capacity: -Logical Unreferenced Percentage: -

cluster1::> storage aggregate show -instance Aggregate: MyFabricPool ... Composite: true Capacity Tier Used Size: ...

<pre>cluster1::> volume show-footprint</pre>			
Vserver : vsl Volume : rootvol			
Feature	Used	Used%	
Volume Footprint	KB	00	
Volume Guarantee	MB	00	
Flexible Volume Metadata	KB	00	
Delayed Frees	KB	00	
Total Footprint	MB	010	
Vserver : vsl Volume : vol			
Feature	Used	Used%	
Volume Footprint	кв	 %	
Footprint in Performance Tier	KB	00	
Footprint in Amazon01	KB	00	
Flexible Volume Metadata	MB	00	
Delayed Frees	KB	00	
Total Footprint	MB	00	

2. Take one of the following actions as needed:

If you want to	Then
Change the tiering policy of a volume	Follow the procedure in Managing storage tiering by modifying a volume's tiering policy or tiering minimum cooling period.
Increase the FabricPool licensed usage limit	Contact your NetApp or partner sales representative. NetApp Support
Increase the storage space of the cloud tier	Contact the provider of the object store that you use for the cloud tier.

Related information

- storage aggregate object
- storage aggregate show

• storage aggregate show-space

Modify an ONTAP volume's tiering policy and minimum cooling period

You can change the tiering policy of a volume to control whether data is moved to the cloud tier when it becomes inactive (*cold*). For a volume with the *snapshot-only* or auto tiering policy, you can also specify the tiering minimum cooling period that user data must remain inactive before it is moved to the cloud tier.

Before you begin

Changing a volume to the auto tiering policy or modifying the tiering minimum cooling period requires ONTAP 9.4 or later.

About this task

Changing the tiering policy of a volume changes only the subsequent tiering behavior for the volume. It does not retroactively move data to the cloud tier.

Changing the tiering policy might affect how long it takes for data to become cold and be moved to the cloud tier.

What happens when you modify the tiering policy of a volume in FabricPool



When in an SVM DR relationship, source and destination volumes do not need to use FabricPool aggregates, but they must use the same tiering policy.

Steps

1. Modify the tiering policy for an existing volume by using the volume modify command with the -tiering-policy parameter:

You can specify one of the following tiering policies:

- ° snapshot-only (default)
- ° auto
- ° all
- ° none

Types of FabricPool tiering policies

2. If the volume uses the snapshot-only or auto tiering policy and you want to modify the tiering minimum cooling period, use the volume modify command with the -tiering-minimum-cooling-days optional parameter in the advanced privilege level.

You can specify a value between 2 and 183 for the tiering minimum cooling days. If you are using a version of ONTAP earlier than 9.8, you can specify a value between 2 and 63 for the tiering minimum cooling days.

Example of modifying the tiering policy and the tiering minimum cooling period of a volume

The following example changes the tiering policy of the volume "myvol" in the SVM "vs1" to auto and the tiering minimum cooling period to 45 days:

```
cluster1::> volume modify -vserver vs1 -volume myvol
-tiering-policy auto -tiering-minimum-cooling-days 45
```

Archive volumes with FabricPool (video)

This video shows a quick overview of using System Manager to archive a volume to a cloud tier with FabricPool.

NetApp video: Archiving volumes with FabricPool (backup + volume move)

Related information

NetApp TechComm TV: FabricPool playlist

Modify an ONTAP volume's default FabricPool tiering policy

You can change a volume's default tiering policy for controlling user data retrieval from the cloud tier to performance tier by using the -cloud-retrieval-policy option introduced in ONTAP 9.8.

Before you begin

- Modifying a volume using the -cloud-retrieval-policy option requires ONTAP 9.8 or later.
- You must have the advanced privilege level to perform this operation.
- You should understand the behavior of tiering policies with -cloud-retrieval-policy.

How tiering policies work with cloud migration

Step

1. Modify the tiering policy behavior for an existing volume by using the volume modify command with the -cloud-retrieval-policy option:

```
volume create -volume <volume_name> -vserver <vserver_name> - tiering-
policy <policy name> -cloud-retrieval-policy
```

vol modify -volume fp_volume4 -vserver vs0 -cloud-retrieval-policy promote

Set thresholds on ONTAP FabricPool per-node put rate

As a storage admin, you can use PUT throttling to set an upper threshold on the maximum per-node put rate.

PUT throttling is useful when network resources or the object store endpoint are resource constrained. Although rare, resource constraints can occur with underpowered object stores or during the first days of FabricPool usage when TB or PB of cold data begins to tier out.

PUT throttling is per node. The minimum PUT throttling put-rate-limit is 8MB/s. Setting the put-rate-limit to a value less than 8MB/s will result in 8MB/s throughput on that node. Multiple nodes, tiering concurrently, might consume more bandwidth and potentially saturate a network link with extremely limited capacity.



FabricPool PUT operations do not compete for resources with other applications. FabricPool PUT operations are automatically placed at a lower priority ("bullied") by client applications and other ONTAP workloads, such as SnapMirror. PUT throttling using put-rate-limit might be useful for reducing network traffic associated with FabricPool tiering, but it is unrelated to concurrent ONTAP traffic.

Before you begin

Advanced privilege level is required.

Steps

1. Throttle FabricPool PUT operations using the ONTAP CLI:

```
storage aggregate object-store put-rate-limit modify -node <name>
-default <true|false> -put-rate-bytes-limit <integer>[KB|MB|GB|TB|PB]
```

Related information

• storage aggregate object-store put-rate-limit modify

Customize ONTAP FabricPool object deletion and defragmentation

FabricPool does not delete blocks from attached object stores. Instead, FabricPool deletes objects after a certain percentage of the blocks in the object are no longer referenced by ONTAP.

For example, there are 1,024 4KB blocks in a 4MB object tiered to Amazon S3. Defragmentation and deletion do not occur until less than 205 4KB blocks (20% of 1,024) are being referenced by ONTAP. When enough (1,024) blocks have zero references, their original 4MB objects are deleted, and a new object is created.

You can customize the unreclaimed space threshold percentage and set it to different default levels for different object stores. The default settings are:

Object Store	ONTAP 9.8 and later	ONTAP 9.7 to 9.4	ONTAP 9.3 and earlier	Cloud Volumes ONTAP
Amazon S3	20%	20%	0%	30%
Google Cloud Storage	20%	12%	n/a	35%
Microsoft Azure Blob Storage	25%	15%	n/a	35%
NetApp ONTAP S3	40%	n/a	n/a	n/a

NetApp	40%	40%	0%	n/a
StorageGRID				

Unreclaimed space threshold

Changing the default unreclaimed space threshold settings will increase or decrease the accepted amount of object fragmentation. Reducing fragmentation will reduce the amount of physical capacity used by the cloud tier at the expense of additional object store resources (reads and writes).

Threshold reduction

To avoid additional expenses, consider reducing the unreclaimed space thresholds when using object store pricing schemes that reduce the cost of storage but increase the cost of reads. Examples include Amazon's Standard-IA and Azure Blob Storage's Cool.

For example, tiering a volume of 10-year-old projects that has been saved for legal reasons might be less expensive when using a pricing scheme such as Standard-IA or Cool than it would be when using standard pricing schemes. Although reads are more expensive for such a volume, including reads required by object defragmentation, they are unlikely to occur frequently.

Threshold increases

Alternatively, consider increasing unreclaimed space thresholds if object fragmentation causes significantly more object store capacity to be used than necessary for the data being referenced by ONTAP. For example, using an unreclaimed space threshold of 20% in a worst-case scenario where all objects are equally fragmented to the maximum allowable extent means that it is possible for 80% of total capacity in the cloud tier to be unreferenced by ONTAP. For example:

2TB referenced by ONTAP + 8TB unreferenced by ONTAP = 10TB total capacity used by the cloud tier.

In this situation, it might be advantageous to increase the unreclaimed space threshold or increase volume minimum cooling days to reduce the capacity used by unreferenced blocks.

As the system defragments objects and increases their storage efficiency, it might fragment the underlying files by writing referenced blocks to new, more efficient objects. If you significantly increase the unreclaimed space threshold, you can create objects that are more storage efficient but have reduced sequential read performance.



This additional activity results in increased costs from third party S3 providers, such as AWS, Azure, and Google.

NetApp recommends avoiding increasing the unreclaimed space threshold above 60%.

Change the unreclaimed space threshold

You can customize the unreclaimed space threshold percentage for different object stores.

Before you begin

Advanced privilege level is required.

Steps

1. To change the default unreclaimed space threshold, customize and run the following command:

```
storage aggregate object-store modify -aggregate <name> -object-store
-name <name> -unreclaimed-space-threshold <%> (0%-99%)
```

Related information

• storage aggregate object-store modify

Promote ONTAP data to the performance tier

Beginning with ONTAP 9.8, if you are a cluster administrator at the advanced privilege level, you can proactively promote data to the performance tier from the cloud tier using a combination of the tiering-policy and the cloud-retrieval-policy setting.

About this task

You might do this if you want to stop using FabricPool on a volume, or if you have a snapshot-only tiering policy and you want to bring restored snapshot data back to the performance tier.

Promote all data from a FabricPool volume to the performance tier

You can proactively retrieve all data on a FabricPool volume in the cloud tier and promote it to the performance tier.

Steps

 Use the volume modify command to set tiering-policy to none and cloud-retrieval-policy to promote.

```
volume modify -vserver <vserver-name> -volume <volume-name> -tiering
-policy none -cloud-retrieval-policy promote
```

Promote file system data to the performance tier

You can proactively retrieve active file system data from a restored snapshot in the cloud tier and promote it to the performance tier.

Steps

1. Use the volume modify command to set tiering-policy to snapshot-only and cloudretrieval-policy to promote.

```
volume modify -vserver <vserver-name> -volume <volume-name> -tiering
-policy snapshot-only cloud-retrieval-policy promote
```

Check the status of a performance tier promotion

You can check the status of performance tier promotion to determine when the operation is complete.

Steps

1. Use the volume object-store command with the tiering option to check the status of the performance tier promotion.

```
volume object-store tiering show [ -instance | -fields <fieldname>, ...
] [ -vserver <vserver name> ] *Vserver
[[-volume] <volume name>] *Volume [ -node <nodename> ] *Node Name [ -vol
-dsid <integer> ] *Volume DSID
[ -aggregate <aggregate name> ] *Aggregate Name
```

```
volume object-store tiering show v1 -instance
                                  Vserver: vsl
                                   Volume: v1
                                Node Name: node1
                              Volume DSID: 1023
                           Aggregate Name: al
                                    State: ready
                      Previous Run Status: completed
                 Aborted Exception Status: -
               Time Scanner Last Finished: Mon Jan 13 20:27:30 2020
                 Scanner Percent Complete: -
                      Scanner Current VBN: -
                         Scanner Max VBNs: -
      Time Waiting Scan will be scheduled: -
                           Tiering Policy: snapshot-only
     Estimated Space Needed for Promotion: -
                        Time Scan Started: -
  Estimated Time Remaining for scan to complete: -
                    Cloud Retrieve Policy: promote
```

Trigger scheduled migration and tiering

Beginning with ONTAP 9.8, you can trigger a tiering scan request at any time when you prefer not to wait for the default tiering scan.

Steps

1. Use the volume object-store command with the trigger option to request migration and tiering.

volume object-store tiering trigger [-vserver <vserver name>] *VServer Name [-volume] <volume name> *Volume Name
Manage FabricPool mirrors

Learn about ONTAP FabricPool mirrors

To ensure data is accessible in data stores in the event of a disaster, and to enable you to replace a data store, you can configure a FabricPool mirror by adding a second data store to synchronously tier data to two data stores. You can add a second data store to new or existing FabricPool configurations, monitor the mirror status, display FabricPool mirror details, promote a mirror, and remove a mirror. You must be running ONTAP 9.7 or later.

Create an ONTAP FabricPool mirror

To create a FabricPool mirror, you attach two object stores to a single FabricPool. You can create a FabricPool mirror either by attaching a second object store to an existing, single object store FabricPool configuration, or you can create a new, single object store FabricPool configuration and then attach a second object store to it. You can also create FabricPool mirrors on MetroCluster configurations.

Before you begin

- You must have already created the two object stores using the storage aggregate object-store config command.
- If you are creating FabricPool mirrors on MetroCluster configurations:
 - You must have already set up and configured the MetroCluster
 - You must have created the object store configurations on the selected cluster.

If you are creating FabricPool mirrors on both clusters in a MetroCluster configuration, you must have created object store configurations on both of the clusters.

- If you are not using on premises object stores for MetroCluster configurations, you should ensure that one of the following scenarios exists:
 - · Object stores are in different availability zones
 - · Object stores are configured to keep copies of objects in multiple availability zones

Setting up object stores for FabricPool in a MetroCluster configuration

About this task

The object store you use for the FabricPool mirror must be different from the primary object store.

The procedure for creating a FabricPool mirror is the same for both MetroCluster and non-MetroCluster configurations.

Steps

1. If you are not using an existing FabricPool configuration, create a new one by attaching an object store to an local tier using the storage aggregate object-store attach command.

This example creates a new FabricPool by attaching an object store to an local tier.

```
cluster1::> storage aggregate object-store attach -aggregate aggr1 -name
my-store-1
```

2. Attach a second object store to the local tier using the storage aggregate object-store mirror command.

This example attaches a second object store to an local tier to create a FabricPool mirror.

```
cluster1::> storage aggregate object-store mirror -aggregate aggr1 -name
my-store-2
```

Related information

- storage aggregate object-store attach
- storage aggregate object-store config
- storage aggregate object-store mirror

Display ONTAP FabricPool mirror details

You can display details about a FabricPool mirror to see what object stores are in the configuration and whether the object store mirror is in sync with the primary object store.

Step

1. Display information about a FabricPool mirror using the storage aggregate object-store show command.

This example displays the details about the primary and mirror object stores in a FabricPool mirror.

```
cluster1::> storage aggregate object-store show
```

Aggregate	Object Store Name	Availability	Mirror Type
aggr1	my-store-1 my-store-2	available available	primary mirror

This example displays details about the FabricPool mirror, including whether the mirror is degraded due to a resync operation.

cluster1::> storage aggregate object-store show -fields mirror-type,ismirror-degraded

aggregate	object-store-name	mirror-type	is-mirror-degraded
aggr1	my-store-1 my-store-2	primary mirror	false

Related information

storage aggregate object-store show

Promote an ONTAP FabricPool mirror

You can reassign the object store mirror as the primary object store by promoting it. When the object store mirror becomes the primary, the original primary automatically becomes the mirror.

Before you begin

- The FabricPool mirror must be in sync
- The object store must be operational

About this task

You can replace the original object store with an object store from a different cloud provider. For instance, your original mirror might be an AWS object store, but you can replace it with an Azure object store.

Steps

1. Verify that the FabricPool mirror is in sync using the storage aggregate object-store showresync-status command. If the FabricPool mirror is in sync, no entries are displayed. If the mirror is not in sync, wait for the resync to complete.

```
aggregate1::> storage aggregate object-store show-resync-status
-aggregate aggr1
```

Aggregate	Primary	Mirror	Complete Percentage
1	1	2	
aggri	my-store-1	my-store-2	40%

2. Promote an object store mirror by using the storage aggregate object-store modify -aggregate command.

cluster1::> storage aggregate object-store modify -aggregate aggr1 -name
my-store-2 -mirror-type primary

Related information

- storage aggregate object-store modify
- storage aggregate object-store show-resync-status

Remove an ONTAP FabricPool mirror

You can remove a FabricPool mirror if you no longer need to replicate an object store.

Before you begin

The primary object store must be operational; otherwise, the command fails.

Step

1. Remove an object store mirror in a FabricPool by using the storage aggregate object-store unmirror -aggregate command.

cluster1::> storage aggregate object-store unmirror -aggregate aggr1

Related information

• storage aggregate object-store unmirror

Replace an existing object store with an ONTAP FabricPool mirror

You can use FabricPool mirror technology to replace one object store with another one. The new object store does not have to use the same cloud provider as the original object store.

About this task

You can replace the original object store with an object store that uses a different cloud provider. For instance, your original object store might use AWS as the cloud provider, but you can replace it with an object store that uses Azure as the cloud provider, and vice versa. However, the new object store must retain the same object size as the original.

Steps

1. Create a FabricPool mirror by adding a new object store to an existing FabricPool using the storage aggregate object-store mirror command.

cluster1::> storage aggregate object-store mirror -aggregate aggr1
-object-store-name my-AZURE-store

 Monitor the mirror resync status using the storage aggregate object-store show-resyncstatus command.

cluster1::> storage aggregate object-store show-resync-status -aggregate
aggr1

			Complete
Aggregate	Primary	Mirror	Percentage
aggrl	my-AWS-store	my-AZURE-store	e 40%

3. Verify the mirror is in sync using the storage aggregate object-store> show -fields mirror-type, is-mirror-degraded command.

```
cluster1::> storage aggregate object-store show -fields mirror-type,is-
mirror-degraded
```

aggregate object-store-name mirror-type is-mirror-degraded aggr1 my-AWS-store primary my-AZURE-store mirror false

 Swap the primary object store with the mirror object store using the storage aggregate objectstore modify command.

cluster1::> storage aggregate object-store modify -aggregate aggr1
-object-store-name my-AZURE-store -mirror-type primary

5. Display details about the FabricPool mirror using the storage aggregate object-store show -fields mirror-type, is-mirror-degraded command.

This example displays the information about the FabricPool mirror, including whether the mirror is degraded (not in sync).

cluster1::> storage aggregate object-store show -fields mirror-type, ismirror-degraded

aggregate	object-store-name	mirror-type	is-mirror-degraded
aggrl	my-AZURE-store	primary	-
	my-AWS-store	mirror	false

6. Remove the FabricPool mirror using the storage aggregate object-store unmirror command.

cluster1::> storage aggregate object-store unmirror -aggregate aggr1

7. Verify that the FabricPool is back in a single object store configuration using the storage aggregate object-store show -fields mirror-type, is-mirror-degraded command.

cluster1::> storage aggregate object-store show -fields mirror-type,ismirror-degraded

```
aggregateobject-store-name mirror-typeis-mirror-degradedaggr1my-AZURE-storeprimary-
```

Related information

- storage aggregate object-store mirror
- storage aggregate object-store modify
- storage aggregate object-store show-resync-status
- · storage aggregate object-store show
- storage aggregate object-store unmirror

Replace a FabricPool mirror in an ONTAP MetroCluster configuration

If one of the object stores in a FabricPool mirror is destroyed or becomes permanently unavailable on a MetroCluster configuration, you can make the object store the mirror if it is not the mirror already, remove the damaged object store from FabricPool mirror, and then add a new object store mirror to the FabricPool.

Steps

1. If the damaged object store is not already the mirror, make the object store the mirror with the storage aggregate object-store modify command.

```
storage aggregate object-store modify -aggregate -aggregate fp_aggr1_A01
-name mcc1_ostore1 -mirror-type mirror
```

2. Remove the object store mirror from the FabricPool by using the storage aggregate object-store unmirror command.

```
storage aggregate object-store unmirror -aggregate <aggregate name>
-name mcc1 ostore1
```

3. You can force tiering to resume on the primary data store after you remove the mirror data store by using the storage aggregate object-store modify with the -force-tiering-on-metrocluster true option.

The absence of a mirror interferes with the replication requirements of a MetroCluster configuration.

```
storage aggregate object-store modify -aggregate <aggregate name> -name
mcc1 ostore1 -force-tiering-on-metrocluster true
```

 Create a replacement object store by using the storage aggregate object-store config create command.

```
storage aggregate object-store config create -object-store-name
mcc1_ostore3 -cluster clusterA -provider-type SGWS -server <SGWS-server-
l> -container-name <SGWS-bucket-l> -access-key <key> -secret-password
<password> -encrypt <true|false> -provider <provider-type> -is-ssl
-enabled <true|false> ipspace <IPSpace>
```

5. Add the object store mirror to the FabricPool mirror using the storage aggregate object-store mirror command.

```
storage aggregate object-store mirror -aggregate aggr1 -name
mcc1 ostore3-mc
```

6. Display the object store information using the storage aggregate object-store show command.

```
storage aggregate object-store show -fields mirror-type, is-mirror-degraded
```

aggregate	object-store-name	mirror-type	is-mirror-degraded
aggrl	mccl_ostorel-mc	primary	-
	mccl_ostore3-mc	mirror	true

 Monitor the mirror resync status using the storage aggregate object-store show-resyncstatus command.

```
storage aggregate object-store show-resync-status -aggregate aggr1
```

			Complete
Aggregate	Primary	Mirror	Percentage
aggrl	mcc1_ostore1-	-mc mcc1_ostore3-	mc 40%

Related information

- storage aggregate object-store config create
- storage aggregate object-store mirror
- storage aggregate object-store modify
- storage aggregate object-store show
- storage aggregate object-store show-resync-status
- storage aggregate object-store unmirror

ONTAP commands for managing FabricPool resources

You use the storage aggregate object-store commands to manage object stores for FabricPool. You use the storage aggregate commands to manage aggregates for FabricPool. You use the volume commands to manage volumes for FabricPool.

If you want to	Use this command:
Define the configuration for an object store so that ONTAP can access it	storage aggregate object-store config create
Modify object store configuration attributes	storage aggregate object-store config modify
Rename an existing object store configuration	storage aggregate object-store config rename
Delete the configuration of an object store	storage aggregate object-store config delete
Display a list of object store configurations	storage aggregate object-store config show
Attach a second object store to a new or existing FabricPool as a mirror	storage aggregate object-store mirror with the -aggregate and -name parameter in the admin privilege level
Remove an object store mirror from an existing FabricPool mirror	storage aggregate object-store unmirror with the -aggregate and -name parameter in the admin privilege level
Monitor FabricPool mirror resync status	storage aggregate object-store show- resync-status
Display FabricPool mirror details	storage aggregate object-store show

Promote an object store mirror to replace a primary object store in a FabricPool mirror configuration	storage aggregate object-store modify with the -aggregate parameter in the admin privilege level
Test the latency and performance of an object store without attaching the object store to an aggregate	storage aggregate object-store profiler start with the -object-store-name and -node parameter in the advanced privilege level
Monitor the object store profiler status	storage aggregate object-store profiler show with the -object-store-name and -node parameter in the advanced privilege level
Abort the object store profiler when it is running	storage aggregate object-store profiler abort with the -object-store-name and -node parameter in the advanced privilege level
Attach an object store to an aggregate for using FabricPool	storage aggregate object-store attach
Attach an object store to an aggregate that contains a FlexGroup volume for using FabricPool	storage aggregate object-store attach with the allow-flexgroup true
Display details of the object stores that are attached to FabricPool-enabled aggregates	storage aggregate object-store show
Display the aggregate fullness threshold used by the tiering scan	storage aggregate object-store show with the -fields tiering-fullness-threshold parameter in the advanced privilege level
Display space utilization of the object stores that are attached to FabricPool-enabled aggregates	storage aggregate object-store show- space
Enable inactive data reporting on an aggregate that is not used for FabricPool	storage aggregate modify with the -is -inactive-data-reporting-enabled true parameter
Display whether inactive data reporting is enabled on an aggregate	storage aggregate show with the -fields is- inactive-data-reporting-enabled parameter
Display information about how much user data is cold within an aggregate	storage aggregate show-space with the -fields performance-tier-inactive-user- data,performance-tier-inactive-user- data-percent parameter

Create a volume for FabricPool, including specifying the following:	volume create
The tiering policy	 You use the -tiering-policy parameter to specify the tiering policy.
 The tiering minimum cooling period (for the snapshot-only or auto tiering policy) 	• You use the -tiering-minimum-cooling -days parameter in the advanced privilege level to specify the tiering minimum cooling period.
Modify a volume for FabricPool, including modifying the following:	volume modify
The tiering policy	 You use the -tiering-policy parameter to specify the tiering policy.
 The tiering minimum cooling period (for the snapshot-only or auto tiering policy) 	• You use the -tiering-minimum-cooling -days parameter in the advanced privilege level to specify the tiering minimum cooling period.
Display FabricPool information related to a volume, including the following:	volume show
 The tiering minimum cooling period How much user data is cold 	• You use the -fields tiering-minimum- cooling-days parameter in the advanced privilege level to display the tiering minimum cooling period.
	• You use the -fields performance-tier- inactive-user-data, performance-tier- inactive-user-data-percent parameter to display how much user data is cold.
Move a volume in to or out of FabricPool	volume move start You use the -tiering -policy optional parameter to specify the tiering policy for the volume.
Modify the threshold for reclaiming unreferenced space (the defragmentation threshold) for FabricPool	storage aggregate object-store modify with the -unreclaimed-space-threshold parameter in the advanced privilege level
Modify the threshold for the percent full the aggregate becomes before the tiering scan begins tiering data for FabricPool	storage aggregate object-store modify with the -tiering-fullness-threshold parameter in the advanced privilege level
FabricPool continues to tier cold data to a cloud tier until the local tier reaches 98% capacity.	
Display the threshold for reclaiming unreferenced space for FabricPool	storage aggregate object-store show or storage aggregate object-store show- space command with the -unreclaimed-space -threshold parameter in the advanced privilege level

Related information

- storage aggregate modify
- storage aggregate object
- storage aggregate show-space

SVM data mobility

SVM data mobility overview

Beginning with ONTAP 9.10.1, cluster administrators can non-disruptively relocate an SVM from a source cluster to a destination cluster to manage capacity and load balancing, or to enable equipment upgrades or data center consolidations by using the ONTAP CLI.

This non-disruptive SVM relocation capability is supported on AFF platforms in ONTAP 9.10.1 and 9.11.1. Beginning with ONTAP 9.12.1, this capability is supported on both FAS and AFF platforms and on hybrid aggregates.

The SVM's name and UUID remain unchanged after migration, as well as the data LIF name, IP address, and object names, such as the volume name. The UUID of the objects in the SVM will be different.

SVM migration workflow

The diagram depicts the typical workflow for an SVM migration. You start an SVM migration from the destination cluster. You can monitor the migration from either the source or the destination. You can perform a manual cutover or an automatic cutover. An automatic cutover is performed by default.



SVM migration platform support

Controller family	ONTAP versions supported
AFF A-series	ONTAP 9.10.1 and later
AFF C-series	ONTAP 9.12.1 patch 4 and later
FAS	ONTAP 9.12.1 and later



When migrating from an AFF cluster to a FAS cluster with hybrid aggregates, auto volume placement will attempt to perform a like to like aggregate match. For example, if the source cluster has 60 volumes, the volume placement will try to find an AFF aggregate on the destination to place the volumes. When there is not sufficient space on the AFF aggregates, the volumes will be placed on aggregates with non-flash disks.

Scalability support by ONTAP version

ONTAP version	HA pairs in source and destination
ONTAP 9.14.1	12
ONTAP 9.13.1	6
ONTAP 9.11.1	3
ONTAP 9.10.1	1

Network infrastructure performance requirements for TCP round trip time (RTT) between the source and the destination cluster

Depending on the ONTAP version installed on the cluster, the network connecting the source and destination clusters must have a maximum round trip time as indicated:

ONTAP version	Maximum RTT
ONTAP 9.12.1 and later	10ms
ONTAP 9.11.1 and earlier	2ms

Maximum supported volumes per SVM

Source	Destination	ONTAP 9.14.1	ONTAP 9.13.1	ONTAP 9.12.1	ONTAP 9.11.1 and earlier
AFF	AFF	400	200	100	100
FAS	FAS	80	80	80	N/A
FAS	AFF	80	80	80	N/A
AFF	FAS	80	80	80	N/A

Prerequisites

Before initiating an SVM migration, you must meet the following prerequisites:

- · You must be a cluster administrator.
- The source and destination clusters must be peered to each other.
- The source and destination clusters must have the SnapMirror synchronous license installed. This license is included with ONTAP One.
- All nodes in the source cluster must be running ONTAP 9.10.1 or later. For specific ONTAP array controller support, see Hardware Universe.
- All nodes in the source cluster must be running the same ONTAP version.
- All nodes in the destination cluster must be running the same ONTAP version.
- The destination cluster ONTAP version must be at the same or no more than two major newer versions as the source cluster.
- The source and destination clusters must support the same IP subnet for data LIF access.
- The source SVM must contain fewer than the maximum number of supported data volumes for the release.
- Sufficient space for volume placement must be available on the destination.
- The Onboard Key Manager or external key management must be configured at the cluster level on the destination if the source SVM has encrypted volumes.
 - In this case, key managers configured at the SVM level on the source will not migrate to the destination. The destination will use the cluster-level key manager.
- If the source has encrypted volumes and is configured for NetApp Aggregate Encryption (NAE), the destination must also be configured for NAE.
- If you are migrating an SVM between a non-MetroCluster configuration and a MetroCluster configuration, or between two MetroCluster configurations, verify that your configuration meets the following requirements:
 - The source and destination MetroCluster clusters are in a "normal" state. This means that they cannot be in switchover mode or in the "waiting-for-switchback" state.
 - The source and destination MetroCluster clusters aren't in the process of an FC-to-IP transition or a hardware refresh.
 - $\,\circ\,$ The source and destination cluster must both be running ONTAP 9.16.1 or later.
 - $\circ\,$ If the source is a MetroCluster cluster, the SVM subtype is "sync-source" (not "sync-destination").

(

If the destination is a MetroCluster cluster, the SVM created on the destination is always "sync-source". If the destination is a non-MetroCluster cluster, the SVM subtype is always "default".

Best practice

When performing an SVM migration, it is a best practice to leave 30% CPU headroom on both the source cluster and the destination cluster to enable the CPU workload to execute.

SVM operations

You should check for operations that can conflict with an SVM migration:

- · No failover operations are in progress
- WAFLIRON cannot be running

- Fingerprint is not in progress
- Vol move, rehost, clone, create, convert or analytics are not running

Supported and unsupported features

The table indicates the ONTAP features supported by SVM data mobility and the ONTAP releases in which support is available.

For information about ONTAP version interoperability between a source and destination in an SVM migration, see Compatible ONTAP versions for SnapMirror relationships.

Feature	Release first supported	Comments
Autonomous Ransomware Protection	ONTAP 9.12.1	
Cloud Volumes ONTAP	Not supported	
External key manager	ONTAP 9.11.1	
FabricPool	Not supported	
Fanout relationship (the migrating source has a SnapMirror source volume with more than one destination)	ONTAP 9.11.1	
FC SAN	Not supported	
Flash Pool	ONTAP 9.12.1	
FlexCache volumes	Not supported	
FlexGroup	Not supported	
IPsec policies	Not supported	
IPv6 LIFs	Not supported	
ISCSI SAN	Not supported	
Job schedule replication	ONTAP 9.11.1	In ONTAP 9.10.1, job schedules are not replicated during migration and must be manually created on the destination. Beginning with ONTAP 9.11.1, job schedules used by the source are replicated automatically during migration.

Load-sharing mirrors	Not supported		
MetroCluster SVMs	ONTAP 9.16.1	Beginning with ONTAP 9.16.1, the following MetroCluster SVM migrations are supported:	
		 Migrating an SVM between a non-MetroClu configuration and a MetroCluster IP configuration 	ister iration
		 Migrating an SVM between two MetroClusto configurations 	er IP
		 Migrating an SVM between a MetroCluster configuration and a MetroCluster IP configuration 	FC iration
		Note: The source and destination cluster must be running ONTAP 9.16.1 or later to support SM migration.	both /M
		The following MetroCluster SVM migrations are supported for all ONTAP versions:	e not
		 Migrating an SVM between two MetroCluster configurations 	er FC
		 Migrating an SVM between a non-MetroClu configuration and a MetroCluster FC config 	ister uration
		See the prerequisites to migrate an SVM in a MetroCluster configuration.	
NetApp Aggregate Encryption (NAE)	ONTAP 9.11.1	NAE volumes must be placed on NAE supportin destination. If no NAE destination is available, migration operation will fail.	ng
NDMP configurations	Not supported		
NetApp Volume Encryption (NVE)	ONTAP 9.10.1	NVE volumes will be migrated as NVE volumes destination.	s on the
NFS and SMB audit logs	ONTAP 9.13.1	For on-premises SVM migration with audit enabled, you should disable a on the source SVM and then perform the migration.	th audit rm
		Before SVM migration:	
		 Audit log redirect must be enabled on the destination cluster. 	
		 The audit log destination path from the sour SVM must be created on the destination clu 	rce uster.
NFS v3, NFS v4.1, and NFS v4.2	ONTAP 9.10.1		

NFS v4.0	ONTAP 9.12.1	
NFSv4.1 with pNFS	ONTAP 9.14.1	
NVMe over Fabric	Not supported	
Onboard key manager (OKM) with Common Criteria mode enabled on source cluster	Not supported	
Qtrees	ONTAP 9.14.1	
Quotas	ONTAP 9.14.1	
S3	Not supported	
SMB protocol	ONTAP 9.12.1	SMB migrations are disruptive and require a client refresh post migration.
SnapMirror cloud relationships	ONTAP 9.12.1	Beginning with ONTAP 9.12.1, when you migrate an on-premises SVM with SnapMirror cloud relationships, the destination cluster must have the SnapMirror cloud license installed, and it must have enough capacity available to support moving the capacity in the volumes that are being mirrored to the cloud.
SnapMirror asynchronous destination	ONTAP 9.12.1	

SnapMirror asynchronous source	ONTAP 9.11.1	 Transfers SnapMirro migration 	can continue as normal on FlexVol or relationships during most of the
		 Any ongo cutover al they cann completes 	ing transfers are canceled during nd new transfers fail during cutover and ot be restarted until the migration s.
		 Schedule during the after the r 	d transfers that were canceled or missed migration are not automatically started nigration completes.
		i	When a SnapMirror source is migrated, ONTAP does not prevent deletion of the volume after migration until the SnapMirror update takes place. This happens because SnapMirror-related information for migrated SnapMirror source volumes is available only after migration is complete, and after the first update takes place.
SMTape settings	Not supported		
SnapLock	Not supported		
SnapMirror active sync	Not supported		
SnapMirror SVM peer relationships	ONTAP 9.12.1		
SnapMirror SVM disaster recovery	Not supported		
SnapMirror synchronous	Not supported		
Snapshots	ONTAP 9.10.1		
Tamperproof snapshot locking	ONTAP 9.14.1	Tamperproof SnapLock. Sr Compliance r	snapshot locking is not equivalent to napLock Enterprise and SnapLock emain unsupported.
Virtual IP LIFs/BGP	Not supported		
Virtual Storage Console 7.0 and later	Not supported		
Volume clones	Not supported		

vStorage	Not supported	Migration is not allowed when vStorage is enabled. To perform a migration, disable the vStorage option, and then reenable it after migration is completed.
		then reenable it after migration is completed.

Supported operations during migration

The following table indicates volume operations supported within the migrating SVM based on migration state:

Volume operation	SVM migration state		
	In progress	Paused	Cutover
Create	Not allowed	Allowed	Not supported
Delete	Not allowed	Allowed	Not supported
File System Analytics disable	Allowed	Allowed	Not supported
File System Analytics enable	Not allowed	Allowed	Not supported
Modify	Allowed	Allowed	Not supported
Offline/Online	Not allowed	Allowed	Not supported
Move/rehost	Not allowed	Allowed	Not supported
Qtree create/modify	Not allowed	Allowed	Not supported
Quota create/modify	Not allowed	Allowed	Not supported
Rename	Not allowed	Allowed	Not supported
Resize	Allowed	Allowed	Not supported
Restrict	Not allowed	Allowed	Not supported
Snapshot attributes modify	Allowed	Allowed	Not supported
Snapshot autodelete modify	Allowed	Allowed	Not supported
Snapshot create	Allowed	Allowed	Not supported
Snapshot delete	Allowed	Allowed	Not supported
Restore file from snapshot	Allowed	Allowed	Not supported

Migrate an SVM

After an SVM migration has completed, clients are cut over to the destination cluster automatically and the unnecessary SVM is removed from the source cluster. Automatic cutover and automatic source cleanup are enabled by default. If necessary, you can disable client auto-cutover to suspend the migration before cutover occurs and you can also disable automatic source SVM cleanup.

• You can use the -auto-cutover false option to suspend the migration when automatic client cutover normally occurs and then manually perform the cutover later.

Manually cutover clients after SVM migration

• You can use the advance privilege -auto-source-cleanup false option to disable the removal of the source SVM after cutover and then trigger source cleanup manually later, after cutover.

Manually remove source SVM after cutover

Migrate an SVM with automatic cutover enabled

By default, clients are cut over to the destination cluster automatically when the migration is complete, and the unnecessary SVM is removed from the source cluster.

Steps

1. From the destination cluster, run the migration prechecks:

```
dest_cluster> vserver migrate start -vserver SVM_name -source-cluster
  cluster name -check-only true
```

2. From the destination cluster, start the SVM migration:

```
dest_cluster> vserver migrate start -vserver SVM_name -source-cluster
cluster name
```

3. Check the migration status:

dest_cluster> vserver migrate show

The status displays migrate-complete when the SVM migration is finished.

Migrate an SVM with automatic client cutover disabled

You can use the -auto-cutover false option to suspend the migration when automatic client cutover normally occurs and then manually perform the cutover later. See Manually cutover clients after SVM migration.

Steps

1. From the destination cluster, run the migration prechecks:

```
dest_cluster> vserver migrate start -vserver SVM_name -source-cluster
cluster_name -check-only true
```

2. From the destination cluster, start the SVM migration:

dest_cluster> vserver migrate start -vserver SVM_name -source-cluster
cluster name -auto-cutover false

3. Check the migration status:

dest_cluster> vserver migrate show The status displays ready-for-cutover when SVM migration completes the asynchronous data transfers, and it is ready for cutover operation.

Migrate an SVM with source cleanup disabled

You can use the advance privilege -auto-source-cleanup false option to disable the removal of the source SVM

after cutover and then trigger source cleanup manually later, after cutover. See Manually remove source SVM.

Steps

1. From the destination cluster, run the migration prechecks:

```
dest_cluster*> vserver migrate start -vserver SVM_name -source-cluster
cluster name -check-only true
```

2. From the destination cluster, start the SVM migration:

```
dest_cluster*> vserver migrate start -vserver SVM_name -source-cluster
cluster name -auto-source-cleanup false
```

3. Check the migration status:

dest_cluster*> vserver migrate show

The status displays ready-for-source-cleanup when SVM migration cutover is complete, and it is ready to remove the SVM on the source cluster.

Monitor volume migration

In addition to monitoring the overall SVM migration with the vserver migrate show command, you can monitor the migration status of the volumes the SVM contains.

Steps

1. Check volume migration status:

dest clust> vserver migrate show-volume

Pause and resume SVM migration

You might want to pause an SVM migration before the migration cutover begins. You can pause an SVM migration using the vserver migrate pause command.

Pause migration

You can pause an SVM migration before client cutover starts by using the <code>vserver migrate pause</code> command.

Some configuration changes are restricted when a migration operation is in progress; however, beginning with ONTAP 9.12.1, you can pause a migration to fix some restricted configurations and for some failed states so that you can fix configuration issues that might have caused the failure. Some of the failed states that you can fix when you pause SVM migration include the following:

- setup-configuration-failed
- migrate-failed

Steps

1. From the destination cluster, pause the migration:

Resume migrations

When you're ready to resume a paused SVM migration or when an SVM migration has failed, you can use the vserver migrate resume command.

Steps

- 1. Do the following from the destination cluster:
 - a. Resume SVM migration:

vserver migrate resume

b. Verify that the SVM migration has resumed, and monitor the progress:

vserver migrate show

Cancel an SVM migration

If you need to cancel an SVM migration before it completes, you can use the vserver migrate abort command. You can cancel an SVM migration only when the operation is in the paused or failed state. You cannot cancel an SVM migration when the status is "cutover-started" or after cutover is complete. You cannot use the abort option when an SVM migration is in progress.

Steps

1. Check the migration status:

dest_cluster> vserver migrate show -vserver <vserver name>

2. Cancel the migration:

dest cluster> vserver migrate abort -vserver <vserver name>

3. Check the progress of the cancel operation:

dest cluster> vserver migrate show

The migration status shows migrate-aborting while the cancel operation is in progress. When the cancel operation completes, the migration status shows nothing.

Manually cut over clients

By default, client cutover to the destination cluster is performed automatically after the

SVM migration reaches "ready-for-cutover" state. If you choose to disable automatic client cutover, you need to perform the client cutover manually.

Steps

1. Manually execute client cutover:

dest cluster> vserver migrate cutover -vserver <vserver name>

2. Check the status of the cutover operation:

dest_cluster> vserver migrate show

Manually remove source SVM after client cutover

If you performed the SVM migration with source cleanup disabled, you can remove the source SVM manually after client cutover is complete.

Steps

1. Verify they status is ready for source cleanup:

dest_cluster> vserver migrate show

2. Clean up the source:

dest_cluster> vserver migrate source-cleanup -vserver <vserver_name>

HA pair management

Learn about HA pair management in ONTAP clusters

Cluster nodes are configured in high-availability (HA) pairs for fault tolerance and nondisruptive operations. If a node fails or if you need to bring a node down for routine maintenance, its partner can take over its storage and continue to serve data from it. The partner gives back storage when the node is brought back on line.

The HA pair controller configuration consists of a pair of matching FAS/AFF storage controllers (local node and partner node). Each of these nodes is connected to the other's disk shelves. When one node in an HA pair encounters an error and stops processing data, its partner detects the failed status of the partner and takes over all data processing from that controller.

Takeover is the process in which a node assumes control of its partner's storage.

Giveback is the process in which the storage is returned to the partner.

By default, takeovers occur automatically in any of the following situations:

- A software or system failure occurs on a node that leads to a panic. The HA pair controllers automatically fail over to their partner node. After the partner has recovered from the panic and booted up, the node automatically performs a giveback, returning the partner to normal operation.
- A system failure occurs on a node, and the node cannot reboot. For example, when a node fails because

of a power loss, HA pair controllers automatically fail over to their partner node and serve data from the surviving storage controller.



If the storage for a node also loses power at the same time, a standard takeover is not possible.

- Heartbeat messages are not received from the node's partner. This could happen if the partner experienced a hardware or software failure (for example, an interconnect failure) that did not result in a panic but still prevented it from functioning correctly.
- You halt one of the nodes without using the -f or -inhibit-takeover true parameter.



In a two-node cluster with cluster HA enabled, halting or rebooting a node using the -inhibit -takeover true parameter causes both nodes to stop serving data unless you first disable cluster HA and then assign epsilon to the node that you want to remain online.

- You reboot one of the nodes without using the -inhibit-takeover true parameter. (The -onboot parameter of the storage failover command is enabled by default.)
- The remote management device (Service Processor) detects failure of the partner node. This is not applicable if you disable hardware-assisted takeover.

You can also manually initiate takeovers with the storage failover takeover command.

Cluster resiliency and diagnostic improvements

Beginning with ONTAP 9.9.1, the following resiliency and diagnostic additions improve cluster operation:

- **Port monitoring and avoidance**: In two-node switchless cluster configurations, the system avoids ports that experience total packet loss (connectivity loss). In ONTAP 9.8.1 and earlier, this functionality was only available in switched configurations.
- Automatic node failover: If a node cannot serve data across its cluster network, that node should not own any disks. Instead its HA partner should take over, if the partner is healthy.
- **Commands to analyze connectivity issues**: Use the following command to display which cluster paths are experiencing packet loss: network interface check cluster-connectivity show

Learn more about network interface check cluster-connectivity show in the ONTAP command reference.

Learn about hardware-assisted takeovers in ONTAP clusters

Enabled by default, the hardware-assisted takeover feature can speed up the takeover process by using a node's remote management device (Service Processor).

When the remote management device detects a failure, it quickly initiates the takeover rather than waiting for ONTAP to recognize that the partner's heartbeat has stopped. If a failure occurs without this feature enabled, the partner waits until it notices that the node is no longer giving a heartbeat, confirms the loss of heartbeat, and then initiates the takeover.

The hardware-assisted takeover feature uses the following process to avoid that wait:

1. The remote management device monitors the local system for certain types of failures.

- 2. If a failure is detected, the remote management device immediately sends an alert to the partner node.
- 3. Upon receiving the alert, the partner initiates takeover.

System events that trigger hardware-assisted takeover

The partner node might generate a takeover depending on the type of alert it receives from the remote management device (Service Processor).

Alert	Takeover initiated upon receipt?	Description
abnormal_reboot	No	An abnormal reboot of the node occurred.
I2_watchdog_reset	Yes	The system watchdog hardware detected an L2 reset. The remote management device detected a lack of response from the system CPU and reset the system.
loss_of_heartbeat	No	The remote management device is no longer receiving the heartbeat message from the node. This alert does not refer to the heartbeat messages between the nodes in the HA pair; it refers to the heartbeat between the node and its local remote management device.
periodic_message	No	A periodic message is sent during a normal hardware- assisted takeover operation.
power_cycle_via_sp	Yes	The remote management device cycled the system power off and on.
power_loss	Yes	A power loss occurred on the node. The remote management device has a power supply that maintains power for a short period after a power loss, allowing it to report the power loss to the partner.
power_off_via_sp	Yes	The remote management device powered off the system.
reset_via_sp	Yes	The remote management device reset the system.
test	No	A test message is sent to verify a hardware-assisted takeover operation.

Related information

Hardware-assisted (HWassist) takeover - Resolution guide

Learn about automatic takeover and giveback in ONTAP clusters

The automatic takeover and giveback operations can work together to reduce and avoid client outages.

By default, if one node in the HA pair panics, reboots, or halts, the partner node automatically takes over and then returns storage when the affected node reboots. The HA pair then resumes a normal operating state.

Automatic takeovers may also occur if one of the nodes become unresponsive.

Automatic giveback occurs by default. If you would rather control giveback impact on clients, you can disable automatic giveback and use the storage failover modify -auto-giveback false -node <node> command. Before performing the automatic giveback (regardless of what triggered it), the partner node waits for a fixed amount of time as controlled by the -delay- seconds parameter of the storage failover modify command. The default delay is 600 seconds.

This process avoids a single, prolonged outage that includes time required for:

- The takeover operation
- The taken-over node to boot up to the point at which it is ready for the giveback
- The giveback operation

If the automatic giveback fails for any of the non-root aggregates, the system automatically makes two additional attempts to complete the giveback.



During the takeover process, the automatic giveback process starts before the partner node is ready for the giveback. When the time limit of the automatic giveback process expires and the partner node is still not ready, the timer restarts. As a result, the time between the partner node being ready and the actual giveback being performed might be shorter than the automatic giveback time.

What happens during takeover

When a node takes over its partner, it continues to serve and update data in the partner's aggregates and volumes.

The following steps occur during the takeover process:

1. If the negotiated takeover is user-initiated, aggregated data is moved from the partner node to the node that is performing the takeover. A brief outage occurs as the current owner of each aggregate (except for the root aggregate) changes over to the takeover node. This outage is briefer than an outage that occurs during a takeover without aggregate relocation.



A negotiated takeover during panic cannot occur in the case of a panic. A takeover can result from a failure not associated with a panic. A failure is experienced when communication is lost between a node and its partner, also called a heartbeat loss. If a takeover occurs because of a failure, the outage might be longer because the partner node needs time to detect the heartbeat loss.

- You can monitor the progress using the storage failover show-takeover command.
- You can avoid the aggregate relocation during this takeover instance by using the -bypass -optimization parameter with the storage failover takeover command.

Aggregates are relocated serially during planned takeover operations to reduce client outage. If aggregate relocation is bypassed, longer client outage occurs during planned takeover events.

- 2. If the user-initiated takeover is a negotiated takeover, the target node gracefully shuts down, followed by takeover of the target node's root aggregate and any aggregates that were not relocated in the first step.
- 3. Data LIFs (logical interfaces) migrate from the target node to the takeover node, or to any other node in the cluster based on LIF failover rules. You can avoid the LIF migration by using the <code>-skip-lif-migration</code> parameter with the storage failover takeover command. In the case of a user-initiated takeover, data LIFs are migrated before storage takeover begins. In the event of a panic or failure, depending upon

your configuration, data LIFs could be migrated with the storage, or after takeover is complete.

4. Existing SMB sessions are disconnected when takeover occurs.



Due to the nature of the SMB protocol, all SMB sessions are disrupted (except for SMB 3.0 sessions connected to shares with the Continuous Availability property set). SMB 1.0 and SMB 2.x sessions cannot reconnect open file handles after a takeover event; therefore, takeover is disruptive and some data loss could occur.

5. SMB 3.0 sessions that are established to shares with the Continuous Availability property enabled can reconnect to the disconnected shares after a takeover event. If your site uses SMB 3.0 connections to Microsoft Hyper-V and the Continuous Availability property is enabled on the associated shares, takeovers are non-disruptive for those sessions.

What happens if a node performing a takeover panics

If the node that is performing the takeover panics within 60 seconds of initiating takeover, the following events occur:

- The node that panicked reboots.
- After it reboots, the node performs self-recovery operations and is no longer in takeover mode.
- · Failover is disabled.
- If the node still owns some of the partner's aggregates, after enabling storage failover, return these aggregates to the partner using the storage failover giveback command.

What happens during giveback

The local node returns ownership to the partner node when issues are resolved, when the partner node boots up, or when giveback is initiated.

The following process takes place in a normal giveback operation. In this discussion, Node A has taken over Node B. Any issues on Node B have been resolved and it is ready to resume serving data.

- 1. Any issues on Node B are resolved and it displays the following message: Waiting for giveback
- 2. The giveback is initiated by the storage failover giveback command or by automatic giveback if the system is configured for it. This initiates the process of returning ownership of Node B's aggregates and volumes from Node A back to Node B.
- 3. Node A returns control of the root aggregate first.
- 4. Node B completes the process of booting up to its normal operating state.
- 5. As soon as Node B reaches the point in the boot process where it can accept the non-root aggregates, Node A returns ownership of the other aggregates, one at a time, until giveback is complete. You can monitor the progress of the giveback by using the storage failover show-giveback command.



The storage failover show-giveback command does not (nor is it intended to) display information about all operations occurring during the storage failover giveback operation. You can use the storage failover show command to display additional details about the current failover status of the node, such as if the node is fully functional, takeover is possible, and giveback is complete.

I/O resumes for each aggregate after giveback is complete for that aggregate, which reduces its overall outage window.

HA policy and its effect on takeover and giveback

ONTAP automatically assigns an HA policy of CFO (controller failover) and SFO (storage failover) to an aggregate. This policy determines how storage failover operations occur for the aggregate and its volumes.

The two options, CFO and SFO, determine the aggregate control sequence ONTAP uses during storage failover and giveback operations.

Although the terms CFO and SFO are sometimes used informally to refer to storage failover (takeover and giveback) operations, they actually represent the HA policy assigned to the aggregates. For example, the terms SFO aggregate or CFO aggregate simply refer to the aggregate's HA policy assignment.

HA policies affect takeover and giveback operations as follows:

- Aggregates created on ONTAP systems (except for the root aggregate containing the root volume) have an HA policy of SFO. Manually initiated takeover is optimized for performance by relocating SFO (non-root) aggregates serially to the partner before takeover. During the giveback process, aggregates are given back serially after the taken-over system boots and the management applications come online, enabling the node to receive its aggregates.
- Because aggregate relocation operations entail reassigning aggregate disk ownership and shifting control from a node to its partner, only aggregates with an HA policy of SFO are eligible for aggregate relocation.
- The root aggregate always has an HA policy of CFO and is given back at the start of the giveback operation. This is necessary to allow the taken-over system to boot. All other aggregates are given back serially after the taken-over system completes the boot process and the management applications come online, enabling the node to receive its aggregates.



Changing the HA policy of an aggregate from SFO to CFO is a Maintenance mode operation. Do not modify this setting unless directed to do so by a customer support representative.

How background updates affect takeover and giveback

Background updates of the disk firmware will affect HA pair takeover, giveback, and aggregate relocation operations differently, depending on how those operations are initiated.

The following list describes how background disk firmware updates affect takeover, giveback, and aggregate relocation:

- If a background disk firmware update occurs on a disk on either node, manually initiated takeover operations are delayed until the disk firmware update finishes on that disk. If the background disk firmware update takes longer than 120 seconds, takeover operations are aborted and must be restarted manually after the disk firmware update finishes. If the takeover was initiated with the -bypass-optimization parameter of the storage failover takeover command set to true, the background disk firmware update occurring on the destination node does not affect the takeover.
- If a background disk firmware update is occurring on a disk on the source (or takeover) node and the takeover was initiated manually with the -options parameter of the storage failover takeover command set to immediate, takeover operations start immediately.
- If a background disk firmware update is occurring on a disk on a node and it panics, takeover of the panicked node begins immediately.
- If a background disk firmware update is occurring on a disk on either node, giveback of data aggregates is delayed until the disk firmware update finishes on that disk.
- If the background disk firmware update takes longer than 120 seconds, giveback operations are aborted

and must be restarted manually after the disk firmware update completes.

• If a background disk firmware update is occurring on a disk on either node, aggregate relocation operations are delayed until the disk firmware update finishes on that disk. If the background disk firmware update takes longer than 120 seconds, aggregate relocation operations are aborted and must be restarted manually after the disk firmware update finishes. If aggregate relocation was initiated with the -override -destination-checks of the storage aggregate relocation node does not affect aggregate relocation.

Related information

• storage aggregate relocation

ONTAP automatic takeover commands

Automatic takeover is enabled by default on all supported NetApp FAS, AFF, and ASA platforms. You might need to change the default behavior and control when automatic takeovers occur when the partner node reboots, panics, or halts.

If you want takeover to occur automatically when the partner node	Use this command
Reboots or halts	storage failover modify -node nodename -onreboot true
Panics	storage failover modify -node nodename -onpanic true

Enable email notification if the takeover capability is disabled

To receive prompt notification if the takeover capability becomes disabled, you should configure your system to enable automatic email notification for the "takeover impossible" EMS messages:

- ha.takeoverImpVersion
- ha.takeoverImpLowMem
- ha.takeoverImpDegraded
- ha.takeoverImpUnsync
- ha.takeoverImpIC
- ha.takeoverImpHotShelf
- ha.takeoverImpNotDef

ONTAP automatic giveback commands

By default, the take-over partner node automatically gives back storage when the off-line node is brought back on line, thus restoring the high-availability pair relationship. In most cases, this is the desired behavior. If you need to disable automatic giveback - for example, if you want to investigate the cause of the takeover before giving back – you need to be aware of the interaction of non-default settings.

If you want to	Use this command
Enable automatic giveback so that giveback occurs as soon as the taken-over node boots, reaches the Waiting for Giveback state, and the Delay before Auto Giveback period has expired. The default setting is true.	storage failover modify -node <i>nodename</i> -auto-giveback true
Disable automatic giveback. The default setting is true.	storage failover modify -node <i>nodename</i> -auto-giveback false
Note: Setting this parameter to false does not disable automatic giveback after takeover on panic; automatic giveback after takeover on panic must be disabled by setting the -auto-giveback-after-panic parameter to false.	
Disable automatic giveback after takeover on panic (this setting is enabled by default).	storage failover modify -node <i>nodename</i> -auto-giveback-after-panic false
Delay automatic giveback for a specified number of seconds (the default is 600). This option determines the minimum time that a node remains in takeover before performing an automatic giveback.	storage failover modify -node <i>nodename</i> -delay-seconds <i>seconds</i>

How variations of the storage failover modify command affect automatic giveback

The operation of automatic giveback depends on how you configure the parameters of the storage failover modify command.

The following table lists the default settings for the storage failover modify command parameters that apply to takeover events not caused by a panic.

Parameter	Default setting
-auto-giveback <i>true</i> <i>false</i>	true
-delay-seconds integer (seconds)	600
-onreboot true false	true

The following table describes how combinations of the -onreboot and -auto-giveback parameters affect automatic giveback for takeover events not caused by a panic.

storage failover modify	Cause of takeover	Does automatic giveback occur?
parameters used		

-onreboot true	reboot command	Yes
-auto-giveback <i>true</i>	halt command, or power cycle operation issued from the Service Processor	Yes
-onreboot true	reboot command	Yes
-auto-giveback <i>false</i>	halt command, or power cycle operation issued from the Service Processor	No
-onreboot <i>false</i> -auto-giveback <i>true</i>	reboot command	N/A In this case, takeover does not occur
	halt command, or power cycle operation issued from the Service Processor	Yes
-onreboot false	reboot command	No
-auto-giveback <i>false</i>	halt command, or power cycle operation issued from the Service Processor	No

The -auto-giveback parameter controls giveback after panic and all other automatic takovers. If the -onreboot parameter is set to true and a takeover occurs due to a reboot, then automatic giveback is always performed, regardless of whether the -auto-giveback parameter is set to true.

The -onreboot parameter applies to reboots and halt commands issued from ONTAP. When the -onreboot parameter is set to false, a takeover does not occur in the case of a node reboot. Therefore, automatic giveback cannot occur, regardless of whether the -auto-giveback parameter is set to true. A client disruption occurs.

The effects of automatic giveback parameter combinations that apply to panic situations.

The following table lists the storage failover modify command parameters that apply to panic situations:

Parameter	Default setting
-onpanic true false	true
-auto-giveback-after-panic <i>true</i> <i>false</i> (Privilege: Advanced)	true
-auto-giveback <i>true</i> <i>false</i>	true

The following table describes how parameter combinations of the storage failover modify command affect automatic giveback in panic situations.

storage failover parameters used	Does automatic giveback occur after panic?
-onpanic true -auto-giveback true -auto-giveback-after-panic true	Yes
-onpanic true -auto-giveback true -auto-giveback-after-panic false	Yes
-onpanic true -auto-giveback false -auto-giveback-after-panic true	Yes
-onpanic true -auto-giveback false -auto-giveback-after-panic false	No
-onpanic false If -onpanic is set to false, takeover/giveback does not occur, regardless of the value set for -auto-giveback or -auto -giveback-after-panic	No

 (\mathbf{i})

(i)

A takeover can result from a failure not associated with a panic. A *failure* is experienced when communication is lost between a node and its partner, also called a *heartbeat loss*. If a takeover occurs because of a failure, giveback is controlled by the -onfailure parameter instead of the -auto-giveback-after-panic parameter.

When a node panics, it sends a panic packet to its partner node. If for any reason the panic packet is not received by the partner node, the panic can be misinterpreted as a failure. Without receipt of the panic packet, the partner node knows only that communication has been lost, and does not know that a panic has occurred. In this case, the partner node processes the loss of communication as a failure instead of a panic, and giveback is controlled by the -onfailure parameter (and not by the -auto-giveback-after-panic parameter).

Learn more about storage failover modify in the ONTAP command reference.

ONTAP manual takeover commands

You can perform a takeover manually when maintenance is required on the partner, and in other similar situations. Depending on the state of the partner, the command you use to perform the takeover varies.

If you want to	Use this command
Take over the partner node	storage failover takeover
Monitor the progress of the takeover as the partner's aggregates are moved to the node doing the takeover	storage failover show-takeover
Display the storage failover status for all nodes in the cluster	storage failover show

Take over the partner node without migrating LIFs	storage failover takeover -skip-lif -migration-before-takeover true
Take over the partner node even if there is a disk mismatch	storage failover takeover -skip-lif -migration-before-takeover true
Take over the partner node even if there is an ONTAP version mismatch	storage failover takeover -option allow -version-mismatch
Note: This option is only used during the nondisruptive ONTAP upgrade process.	
Take over the partner node without performing aggregate relocation	storage failover takeover -bypass -optimization true
Take over the partner node before the partner has	

Before you issue the storage failover command with the immediate option, you must migrate the data LIFs to another node by using the following command: network interface migrateall -node *node*

Learn more about network interface migrate-all in the ONTAP command reference.

If you specify the storage failover takeover -option immediate command without first migrating the data LIFs, data LIF migration from the node is significantly delayed even if the skip-lif-migration-before-takeover option is not specified.

Similarly, if you specify the immediate option, negotiated takeover optimization is bypassed even if the bypass-optimization option is set to *false*.

Moving epsilon for certain manually initiated takeovers

You should move epsilon if you expect that any manually initiated takeovers could result in your storage system being one unexpected node failure away from a cluster-wide loss of quorum.

About this task

To perform planned maintenance, you must take over one of the nodes in an HA pair. Cluster-wide quorum must be maintained to prevent unplanned client data disruptions for the remaining nodes. In some instances, performing the takeover can result in a cluster that is one unexpected node failure away from cluster-wide loss of quorum.

This can occur if the node being taken over holds epsilon or if the node with epsilon is not healthy. To maintain a more resilient cluster, you can transfer epsilon to a healthy node that is not being taken over. Typically, this would be the HA partner.

Only healthy and eligible nodes participate in quorum voting. To maintain cluster-wide quorum, more than N/2 votes are required (where *N* represents the sum of healthy, eligible, online nodes). In clusters with an even number of online nodes, epsilon adds additional voting weight toward maintaining quorum for the node to which it is assigned.

(i)

Although cluster formation voting can be modified by using the cluster modify -eligibility false command, you should avoid this except for situations such as restoring the node configuration or prolonged node maintenance. If you set a node as ineligible, it stops serving SAN data until the node is reset to eligible and rebooted. NAS data access to the node might also be affected when the node is ineligible.

Steps

- 1. Verify the cluster state and confirm that epsilon is held by a healthy node that is not being taken over:
 - a. Change to the advanced privilege level, confirming that you want to continue when the advanced mode prompt appears (*>):

set -privilege advanced

b. Determine which node holds epsilon:

cluster show

In the following example, Node1 holds epsilon:

Node	Health	Eligibility	Epsilon
Node1	true	true	true
Node2	true	true	false

If the node you want to take over does not hold epsilon, proceed to Step 4.

Learn more about cluster show in the ONTAP command reference.

2. Remove epsilon from the node that you want to take over:

cluster modify -node Nodel -epsilon false

3. Assign epsilon to the partner node (in this example, Node2):

cluster modify -node Node2 -epsilon true

4. Perform the takeover operation:

storage failover takeover -ofnode node_name

5. Return to the admin privilege level:

set -privilege admin

ONTAP manual giveback commands

You can perform a normal giveback, a giveback in which you terminate processes on the partner node, or a forced giveback.



Prior to performing a giveback, you must remove the failed drives in the taken-over system as described in Disks and aggregates management.

If giveback is interrupted

If the takeover node experiences a failure or a power outage during the giveback process, that process stops and the takeover node returns to takeover mode until the failure is repaired or the power is restored.

However, this depends upon the stage of giveback in which the failure occurred. If the node encountered failure or a power outage during partial giveback state (after it has given back the root aggregate), it will not return to takeover mode. Instead, the node returns to partial-giveback mode. If this occurs, complete the process by repeating the giveback operation.

If giveback is vetoed

If giveback is vetoed, you must check the EMS messages to determine the cause. Depending on the reason or reasons, you can decide whether you can safely override the vetoes.

The storage failover show-giveback command displays the giveback progress and shows which subsystem vetoed the giveback, if any. Soft vetoes can be overridden, while hard vetoes cannot be, even if forced. The following tables summarize the soft vetoes that should not be overridden, along with recommended workarounds.

You can review the EMS details for any giveback vetoes by using the following command:

event log show -node * -event gb*

Learn more about event log show in the ONTAP command reference.

Giveback of the root aggregate

These vetoes do not apply to aggregate relocation operations:

Vetoing subsystem module	Workaround
vfiler_low_level	Terminate the SMB sessions causing the veto, or shutdown the SMB application that established the open sessions.
	Overriding this veto might cause the application using SMB to disconnect abruptly and lose data.
Disk Check	All failed or bypassed disks should be removed before attempting giveback. If disks are sanitizing, you should wait until the operation completes.
	Overriding this veto might cause an outage caused by aggregates or volumes going offline due to reservation conflicts or inaccessible disks.

Giveback of the SFO aggregates

These vetoes do not apply to aggregate relocation operations:

Vetoing subsystem module	Workaround
Lock Manager	Gracefully shutdown the SMB applications that have open files, or move those volumes to a different aggregate. Overriding this veto results in loss of SMB lock state, causing disruption and data loss.
Lock Manager NDO	Wait until the locks are mirrored.
	Overriding this veto causes disruption to Microsoft Hyper-V virtual machines.
RAID	Check the EMS messages to determine the cause of the veto:
	If the veto is due to nvfile, bring the offline volumes and aggregates online.
	If disk add or disk ownership reassignment operations are in progress, wait until they complete.
	If the veto is due to an aggregate name or UUID conflict, troubleshoot and resolve the issue.
	If the veto is due to mirror resync, mirror verify, or offline disks, the veto can be overridden and the operation restarts after giveback.
Disk Inventory	Troubleshoot to identify and resolve the cause of the problem.
	The destination node might be unable to see disks belonging to an aggregate being migrated.
	Inaccessible disks can result in inaccessible aggregates or volumes.
Volume Move Operation	Troubleshoot to identify and resolve the cause of the problem.
	This veto prevents the volume move operation from aborting during the important cutover phase. If the job is aborted during cutover, the volume might become inaccessible.

Commands for performing a manual giveback

You can manually initiate a giveback on a node in an HA pair to return storage to the original owner after completing maintenance or resolving any issues that caused the takeover.

If you want to	Use this command
Give back storage to a partner node	storage failover giveback -ofnode nodename

Give back storage even if the partner is not in the waiting for giveback mode	<pre>storage failover giveback -ofnode nodename -require-partner-waiting false Do not use this option unless a longer client outage is acceptable.</pre>
Give back storage even if processes are vetoing the giveback operation (force the giveback)	storage failover giveback -ofnode nodename -override-vetoes true Use of this option can potentially lead to longer client outage, or aggregates and volumes not coming online after the giveback.
Give back only the CFO aggregates (the root aggregate)	storage failover giveback -ofnode <i>nodename</i> -only-cfo-aggregates true
Monitor the progress of giveback after you issue the giveback command	storage failover show-giveback

Testing takeover and giveback in ONTAP clusters

After you configure all aspects of your HA pair, you need to verify that it is operating as expected in maintaining uninterrupted access to both nodes' storage during takeover and giveback operations. Throughout the takeover process, the local (or takeover) node should continue serving the data normally provided by the partner node. During giveback, control and delivery of the partner's storage should return to the partner node.

Steps

- 1. Check the cabling on the HA interconnect cables to make sure that they are secure.
- 2. Verify that you can create and retrieve files on both nodes for each licensed protocol.
- 3. Enter the following command:

storage failover takeover -ofnode partnernode

Learn more about storage failover takeover in the ONTAP command reference.

4. Enter either of the following commands to confirm that takeover occurred:

storage failover show-takeover

storage failover show

If you have the storage failover command's -auto-giveback option enabled:
Node	Partner	Takeover Possible	State Description
node 1	node 2	-	Waiting for giveback
node 2	node 1	false	In takeover, Auto giveback will be initiated in number of seconds

If you have the storage failover command's -auto-giveback option disabled:

Node	Partner	Takeover Possible	State Description
node 1	node 2	-	Waiting for giveback
node 2	node 1	false	In takeover

5. Display all the disks that belong to the partner node (Node2) that the takeover node (Node1) can detect:

storage disk show -home node2 -ownership

The following command displays all disks belonging to Node2 that Node1 can detect: cluster::> storage disk show -home node2 -ownership

Disk	Aggrega te	Home	Owner	DR Home	Home ID	Owner ID	DR Home ID	Reserve r	Pool
1.0.2	-	node2	node2	-	4078312 453	4078312 453	-	4078312 452	Pool0
1.0.3	-	node2	node2	-	4078312 453	4078312 453	-	4078312 452	Pool0

6. Cconfirm that the takeover node (Node1) controls the partner node's (Node2) aggregates:

aggr show -fields home-id, home-name, is-home

aggregate	home-id	home-nameh	is-home
aggr0_1	2014942045	node1	true
aggr0_2	4078312453	node2	false
aggr1_1	2014942045	node1	true
aggr1_2	4078312453	node2	false

During takeover, the "is-home" value of the partner node's aggregates is false.

7. Give back the partner node's data service after it displays the "Waiting for giveback" message:

storage failover giveback -ofnode partnernode

8. Enter either of the following commands to observe the progress of the giveback operation:

storage failover show-giveback

storage failover show

9. Proceed, depending on whether you saw the message that giveback was completed successfully:

If takeover and giveback	Then
Are completed successfully	Repeat Step 2 through Step 8 on the partner node.
Fail	Correct the takeover or giveback failure and then repeat this procedure.

ONTAP commands for monitoring an HA pair

You can use ONTAP commands to monitor the status of the HA pair. If a takeover occurs, you can also determine what caused the takeover.

If you want to check	Use this command
Whether failover is enabled or has occurred, or reasons why failover is not currently possible	storage failover show
View the nodes on which the storage failover HA- mode setting is enabled You must set the value to ha for the node to participate in a storage failover (HA pair) configuration.	storage failover show -fields mode
Whether hardware-assisted takeover is enabled	storage failover hwassist show
The history of hardware-assisted takeover events that have occurred	storage failover hwassist stats show
The progress of a takeover operation as the partner's aggregates are moved to the node doing the takeover	storage failover show-takeover
The progress of a giveback operation in returning aggregates to the partner node	storage failover show-giveback
Whether an aggregate is home during takeover or giveback operations	aggregate show -fields home-id,owner -id,home-name,owner-name,is-home
Whether cluster HA is enabled (applies only to two node clusters)	cluster ha show
The HA state of the components of an HA pair (on systems that use the HA state)	ha-config show This is a Maintenance mode command.

Node states displayed by storage failover show-type commands

The following list describes the node states that the storage failover show command displays.

Node State	Description
------------	-------------

Connected to partner_name, Automatic takeover disabled.	The HA interconnect is active and can transmit data to the partner node. Automatic takeover of the partner is disabled.
Waiting for partner_name, Giveback of partner spare disks pending.	 The local node cannot exchange information with the partner node over the HA interconnect. Giveback of SFO aggregates to the partner is done, but partner spare disks are still owned by the local node. Run the storage failover show-giveback command for more information.
Waiting for partner_name. Waiting for partner lock synchronization.	The local node cannot exchange information with the partner node over the HA interconnect, and is waiting for partner lock synchronization to occur.
Waiting for partner_name. Waiting for cluster applications to come online on the local node.	The local node cannot exchange information with the partner node over the HA interconnect, and is waiting for cluster applications to come online.
Takeover scheduled. target node relocating its SFO aggregates in preparation of takeover.	Takeover processing has started. The target node is relocating ownership of its SFO aggregates in preparation for takeover.
Takeover scheduled. target node has relocated its SFO aggregates in preparation of takeover.	Takeover processing has started. The target node has relocated ownership of its SFO aggregates in preparation for takeover.
Takeover scheduled. Waiting to disable background disk firmware updates on local node. A firmware update is in progress on the node.	Takeover processing has started. The system is waiting for background disk firmware update operations on the local node to complete.
Relocating SFO aggregates to taking over node in preparation of takeover.	The local node is relocating ownership of its SFO aggregates to the taking-over node in preparation for takeover.
Relocated SFO aggregates to taking over node. Waiting for taking over node to takeover.	Relocation of ownership of SFO aggregates from the local node to the taking-over node has completed. The system is waiting for takeover by the taking-over node.
Relocating SFO aggregates to partner_name. Waiting to disable background disk firmware updates on the local node. A firmware update is in progress on the node.	Relocation of ownership of SFO aggregates from the local node to the taking-over node is in progress. The system is waiting for background disk firmware update operations on the local node to complete.

Relocating SFO aggregates to partner_name. Waiting to disable background disk firmware updates on partner_name. A firmware update is in progress on the node.	Relocation of ownership of SFO aggregates from the local node to the taking-over node is in progress. The system is waiting for background disk firmware update operations on the partner node to complete.
Connected to partner_name. Previous takeover attempt was aborted because reason. Local node owns some of partner's SFO aggregates. Reissue a takeover of the partner with the -bypass- optimization parameter set to true to takeover remaining aggregates, or issue a giveback of the partner to return the relocated aggregates.	 The HA interconnect is active and can transmit data to the partner node. The previous takeover attempt was aborted because of the reason displayed under reason. The local node owns some of its partner's SFO aggregates. Either reissue a takeover of the partner node, setting the -bypass-optimization parameter to true to takeover the remaining SFO aggregates, or perform a giveback of the partner to return relocated aggregates.
Connected to partner_name. Previous takeover attempt was aborted. Local node owns some of partner's SFO aggregates. Reissue a takeover of the partner with the -bypass- optimization parameter set to true to takeover remaining aggregates, or issue a giveback of the partner to return the relocated aggregates.	 The HA interconnect is active and can transmit data to the partner node. The previous takeover attempt was aborted. The local node owns some of its partner's SFO aggregates. Either reissue a takeover of the partner node, setting the -bypass-optimization parameter to true to takeover the remaining SFO aggregates, or perform a giveback of the partner to return relocated aggregates.
Waiting for partner_name. Previous takeover attempt was aborted because reason. Local node owns some of partner's SFO aggregates. Reissue a takeover of the partner with the "-bypass- optimization" parameter set to true to takeover remaining aggregates, or issue a giveback of the partner to return the relocated aggregates.	 The local node cannot exchange information with the partner node over the HA interconnect. The previous takeover attempt was aborted because of the reason displayed under reason. The local node owns some of its partner's SFO aggregates. Either reissue a takeover of the partner node, setting the -bypass-optimization parameter to true to takeover the remaining SFO aggregates, or perform a giveback of the partner to return relocated aggregates.
Waiting for partner_name. Previous takeover attempt was aborted. Local node owns some of partner's SFO aggregates. Reissue a takeover of the partner with the "-bypass- optimization" parameter set to true to takeover remaining aggregates, or issue a giveback of the partner to return the relocated aggregates.	 The local node cannot exchange information with the partner node over the HA interconnect. The previous takeover attempt was aborted. The local node owns some of its partner's SFO aggregates. Either reissue a takeover of the partner node, setting the -bypass-optimization parameter to true to takeover the remaining SFO aggregates, or perform a giveback of the partner to return relocated aggregates.

Connected to partner_name. Previous takeover attempt was aborted because failed to disable background disk firmware update (BDFU) on local node.	The HA interconnect is active and can transmit data to the partner node. The previous takeover attempt was aborted because the background disk firmware update on the local node was not disabled.
Connected to partner_name. Previous takeover attempt was aborted because reason.	The HA interconnect is active and can transmit data to the partner node. The previous takeover attempt was aborted because of the reason displayed under reason.
Waiting for partner_name. Previous takeover attempt was aborted because reason.	The local node cannot exchange information with the partner node over the HA interconnect. The previous takeover attempt was aborted because of the reason displayed under reason.
Connected to partner_name. Previous takeover attempt by partner_name was aborted because reason.	The HA interconnect is active and can transmit data to the partner node. The previous takeover attempt by the partner node was aborted because of the reason displayed under reason.
Connected to partner_name. Previous takeover attempt by partner_name was aborted.	The HA interconnect is active and can transmit data to the partner node. The previous takeover attempt by the partner node was aborted.
Waiting for partner_name. Previous takeover attempt by partner_name was aborted because reason.	The local node cannot exchange information with the partner node over the HA interconnect. The previous takeover attempt by the partner node was aborted because of the reason displayed under reason.
Previous giveback failed in module: module name. Auto giveback will be initiated in number of seconds seconds.	The previous giveback attempt failed in module module_name. Auto giveback will be initiated in number of seconds seconds.
	command for more information.
Node owns partner's aggregates as part of the non- disruptive controller upgrade procedure.	The node owns its partner's aggregates due to the non- disruptive controller upgrade procedure currently in progress.
Connected to partner_name. Node owns aggregates belonging to another node in the cluster.	The HA interconnect is active and can transmit data to the partner node. The node owns aggregates belonging to another node in the cluster.
Connected to partner_name. Waiting for partner lock synchronization.	The HA interconnect is active and can transmit data to the partner node. The system is waiting for partner lock synchronization to complete.

Connected to partner_name. Waiting for cluster applications to come online on the local node.	The HA interconnect is active and can transmit data to the partner node. The system is waiting for cluster applications to come online on the local node.
Non-HA mode, reboot to use full NVRAM.	Storage failover is not possible. The HA mode option is configured as non_ha.You must reboot the node to use all of its NVRAM.
Non-HA mode. Reboot node to activate HA.	Storage failover is not possible.The node must be rebooted to enable HA capability.
Non-HA mode.	 Storage failover is not possible. The HA mode option is configured as non_ha. You must run the storage failover modify -mode ha -node nodename command on both nodes in the HA pair and then reboot the nodes to enable HA capability.

Related information

- ONTAP command reference
- cluster ha show

 (\mathbf{i})

ONTAP commands for enabling and disabling storage failover

Use the following commands to enable and disable storage failover functionality.

If you want to	Use this command
Enable takeover	storage failover modify -enabled true -node <i>nodename</i>
Disable takeover	storage failover modify -enabled false -node <i>nodename</i>

You should only disable storage failover if required as part of a maintenance procedure.

Halt or reboot ONTAP nodes without initiating takeover in two-node clusters

You halt or reboot a node in a two-node cluster without initiating takeover when you perform certain hardware maintenance on a node or a shelf and you want to limit down time by keeping the partner node up, or when there are issues preventing a manual takeover and you want to keep the partner node's aggregates up and serving data. Additionally, if technical support is assisting you with troubleshooting problems, they might have you perform this procedure as part of those efforts.

About this task

- Before you inhibit takeover (using the -inhibit-takeover true parameter), you disable cluster HA.
 - In a two-node cluster, cluster HA ensures that the failure of one node does not disable the cluster. However, if you do not disable cluster HA before using the -inhibit-takeover true parameter, both nodes stop serving data.
 - If you attempt to halt or reboot a node before disabling cluster HA, ONTAP issues a warning and instructs you to disable cluster HA.
- You migrate LIFs (logical interfaces) to the partner node that you want to remain online.
- If on the node you are halting or rebooting there are aggregates you want to keep, you move them to the node that you want to remain online.

Steps

- 1. Verify both nodes are healthy:
 - cluster show

For both nodes, true appears in the Health column.

```
cluster::> cluster show
Node Health Eligibility
node1 true true
node2 true true
```

Learn more about cluster show in the ONTAP command reference.

 Migrate all LIFs from the node that you will halt or reboot to the partner node: network interface migrate-all -node node name

Learn more about network interface migrate-all in the ONTAP command reference.

- 3. If on the node you will halt or reboot there are aggregates you want to keep online when the node is down, relocate them to the partner node; otherwise, go to the next step.
 - a. Show the aggregates on the node you will halt or reboot: storage aggregates show -node *node_name*

For example, node1 is the node that will be halted or rebooted:

cluster::> storage aggregates show -node node1 Aggregate Size Available Used% State #Vols Nodes RAID Status _____ _____ ____ _____ aggr0 node 1 0 744.9GB 32.68GB 96% online 2 nodel raid dp, normal aggr1 2.91TB 2.62TB 10% online 8 nodel raid dp, normal aggr2 12 nodel 4.36TB 3.74TB 14% online raid dp, normal test2 aggr 2.18TB 2.18TB 0% online 7 nodel raid dp, normal 4 entries were displayed.

b. Move the aggregates to the partner node:

storage aggregate relocation start -node node_name -destination node_name
-aggregate-list aggregate_name

For example, aggregates aggr1, aggr2 and test2_aggr are being moved from node1 to node2:

storage aggregate relocation start -node node1 -destination node2 -aggregate
-list aggr1,aggr2,test2_aggr

4. Disable cluster HA:

cluster ha modify -configured false

The return output confirms HA is disabled: Notice: HA is disabled



This operation does not disable storage failover.

- 5. Halt or reboot and inhibit takeover of the target node, by using the appropriate command:
 - ° system node halt -node node name -inhibit-takeover true
 - ° system node reboot -node node_name -inhibit-takeover true



In the command output, you will see a warning asking you if you want to proceed, enter y.

6. Verify that the node that is still online is in a healthy state (while the partner is down): cluster show

For the online node, true appears in the Health column.



In the command output, you will see a warning that cluster HA is not configured. You can ignore the warning at this time.

- 7. Perform the actions that required you to halt or reboot the node.
- 8. Boot the offlined node from the LOADER prompt: boot_ontap
- 9. Verify both nodes are healthy: cluster show

For both nodes, true appears in the Health column.



In the command output, you will see a warning that cluster HA is not configured. You can ignore the warning at this time.

10. Reenable cluster HA:

cluster ha modify -configured true

11. If earlier in this procedure you relocated aggregates to the partner node, move them back to their home node; otherwise, go to the next step:

```
storage aggregate relocation start -node node_name -destination node_name
-aggregate-list aggregate name
```

For example, aggregates aggr1, aggr2 and test2_aggr are being moved from node node2 to node node1: storage aggregate relocation start -node *node2* -destination *node1* -aggregate -list *aggr1, aggr2, test2 aggr*

12. Revert LIFs to their home ports:

a. View LIFs that are not at home: network interface show -is-home false

Learn more about network interface show in the ONTAP command reference.

- b. If there are non-home LIFs that were not migrated from the down node, verify it is safe to move them before reverting.
- c. If it is safe to do so, revert all LIFs home. network interface revert *

Learn more about network interface revert in the ONTAP command reference.

Related information

- · cluster ha modify
- storage aggregate relocation start

REST API management with System Manager

REST API management with System Manager

The REST API log captures the API calls that System Manager issues to ONTAP. You can use the log to understand the nature and sequence of the calls needed to perform the various ONTAP administrative tasks.

How System Manager uses the REST API and API log

There are several ways that REST API calls are issued by System Manager to ONTAP.

When does System Manager issue API calls

Here are the most important examples of when System Manager issues ONTAP REST API calls.

Automatic page refresh

System Manager automatically issues API calls in the background to refresh the displayed information, such as on the dashboard page.

Display action by user

One or more API calls are issued when you display a specific storage resource or a collection of resources from the System Manager UI.

Update action by user

An API call is issued when you add, modify, or delete an ONTAP resource from the System Manager UI.

Reissuing an API call

You can also manually reissue an API call by clicking a log entry. This displays the raw JSON output from the call.

More information

ONTAP 9 Automation docs

Accessing the REST API log

You can access the log containing a record of the ONTAP REST API calls made by System Manager. When displaying the log, you can also reissue API calls and review the output.

Steps

1.

At the top of the page, click 🚺 to display the REST API log.

The most recent entries are displayed at the bottom of the page.

- 2. On the left, click **DASHBOARD** and observe the new entries being created for the API calls issued to refresh the page.
- 3. Click **STORAGE** and then click **Qtrees**.

This causes System Manager to issue a specific API call to retrieve a list of the Qtrees.

4. Locate the log entry describing the API call which has the form:

GET /api/storage/qtrees

You will see additional HTTP query parameters included with the entry, such as max_records.

5. Click the log entry to reissue the GET API call and display the raw JSON output.

Example

```
1 {
 2
     "records": [
 3
       {
 4
         "svm": {
 5
           "uuid": "19507946-e801-11e9-b984-00a0986ab770",
 6
           "name": "SMQA",
 7
           " links": {
 8
             "self": {
 9
               "href": "/api/svm/svms/19507946-e801-11e9-b984-
   00a0986ab770"
10
             }
11
           }
         },
12
13
         "volume": {
14
           "uuid": "1e173258-f98b-11e9-8f05-00a0986abd71",
15
          "name": "vol vol test2 dest dest",
           " links": {
16
            "self": {
17
               "href": "/api/storage/volumes/1e173258-f98b-11e9-8f05-
18
   00a0986abd71"
19
            }
20
           }
21
         },
22
         "id": 1,
23
         "name": "test2",
24
         "security style": "mixed",
         "unix permissions": 777,
25
         "export policy": {
26
27
           "name": "default",
           "id": 12884901889,
28
29
           " links": {
30
             "self": {
31
               "href": "/api/protocols/nfs/export-policies/12884901889"
32
             }
33
           }
```

```
34
        },
        "path": "/vol_vol_test2_dest_dest/test2",
35
36
        " links": {
         "self": {
37
38
          "href": "/api/storage/qtrees/1e173258-f98b-11e9-8f05-
00a0986abd71/1"
      }
39
40
       }
41
      },
42
     ],
    "num records": 1,
43
     " links": {
44
       "self": {
45
         "href":
46
 "/api/storage/qtrees?max_records=20&fields=*&name=!%22%22"
47
       }
48 }
49 }
```

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. 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. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at http://www.netapp.com/TM are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.