



Provisioning storage

Cloud Manager 3.6

NetApp
June 10, 2024

Table of Contents

- Provisioning storage 1
 - Provisioning storage 1
 - Tiering inactive data to low-cost object storage 4
 - Using Cloud Volumes ONTAP as persistent storage for Kubernetes 7
 - Encrypting volumes with NetApp Volume Encryption 10
 - Managing existing storage 11
 - Provisioning NFS volumes from the Volume View 17

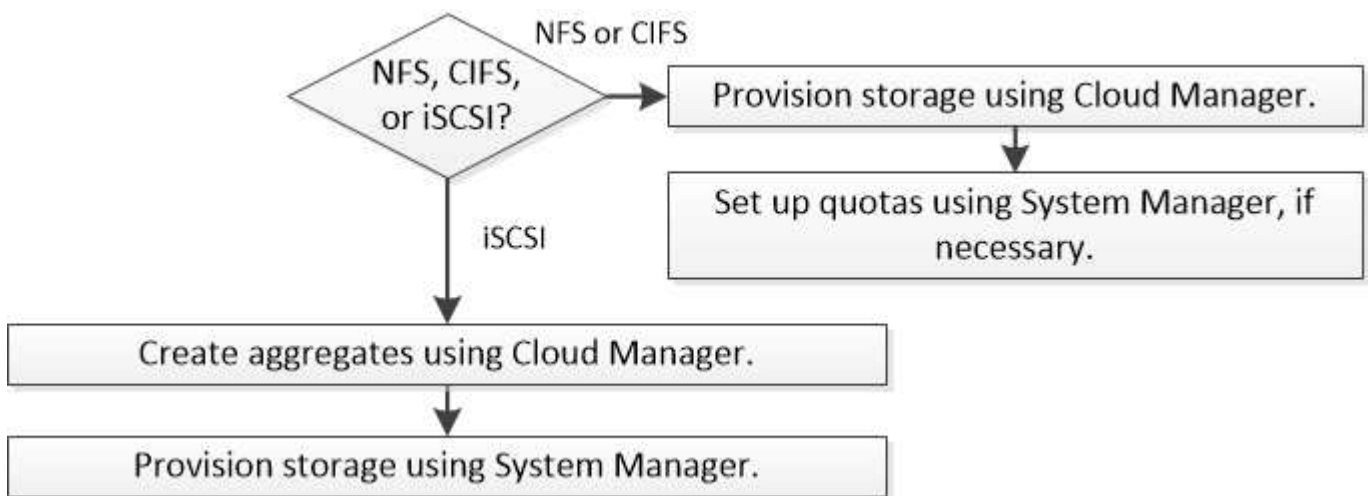
Provisioning storage

Provisioning storage

You can provision additional NFS and CIFS storage for your Cloud Volumes ONTAP systems from Cloud Manager by managing volumes and aggregates. If you need to create iSCSI storage, you should do so from System Manager.



All disks and aggregates must be created and deleted directly from Cloud Manager. You should not perform these actions from another management tool. Doing so can impact system stability, hamper the ability to add disks in the future, and potentially generate redundant cloud provider fees.



Provisioning volumes

If you need more storage after you launch a Cloud Volumes ONTAP system, you can provision new NFS and CIFS volumes from Cloud Manager.

Before you begin

If you want to use CIFS in AWS, you must have set up DNS and Active Directory. For details, see [Networking requirements for Cloud Volumes ONTAP for AWS](#).

Steps

1. On the Working Environments page, double-click the name of the Cloud Volumes ONTAP system on which you want to provision volumes.
2. Create a new volume on any aggregate or on a specific aggregate:

Action	Steps
Create a new volume and let Cloud Manager choose the containing aggregate	Click Add New Volume .

Action	Steps
Create a new volume on a specific aggregate	<ol style="list-style-type: none"> Click the menu icon, and then click Advanced > Advanced allocation. Click the menu for an aggregate. Click Create volume.

- Enter details for the new volume, and then click **Continue**.

Some of the fields in this page are self-explanatory. The following table describes fields for which you might need guidance:

Field	Description
Size	The maximum size that you can enter largely depends on whether you enable thin provisioning, which enables you to create a volume that is bigger than the physical storage currently available to it.
Access control (for NFS only)	An export policy defines the clients in the subnet that can access the volume. By default, Cloud Manager enters a value that provides access to all instances in the subnet.
Permissions and Users / Groups (for CIFS only)	These fields enable you to control the level of access to a share for users and groups (also called access control lists or ACLs). You can specify local or domain Windows users or groups, or UNIX users or groups. If you specify a domain Windows user name, you must include the user's domain using the format domain\username.
Snapshot Policy	A Snapshot copy policy specifies the frequency and number of automatically created NetApp Snapshot copies. A NetApp Snapshot copy is a point-in-time file system image that has no performance impact and requires minimal storage. You can choose the default policy or none. You might choose none for transient data: for example, tempdb for Microsoft SQL Server.

- If you chose the CIFS protocol and the CIFS server has not been set up, specify details for the server in the Create a CIFS Server dialog box, and then click **Save and continue**:

Field	Description
DNS Primary and Secondary IP Address	<p>The IP addresses of the DNS servers that provide name resolution for the CIFS server.</p> <p>The listed DNS servers must contain the service location records (SRV) needed to locate the Active Directory LDAP servers and domain controllers for the domain that the CIFS server will join.</p>
Active Directory Domain to join	The FQDN of the Active Directory (AD) domain that you want the CIFS server to join.
Credentials authorized to join the domain	The name and password of a Windows account with sufficient privileges to add computers to the specified Organizational Unit (OU) within the AD domain.
CIFS server NetBIOS name	A CIFS server name that is unique in the AD domain.

Field	Description
Organizational Unit	The organizational unit within the AD domain to associate with the CIFS server. The default is CN=Computers.
DNS Domain	The DNS domain for the Cloud Volumes ONTAP storage virtual machine (SVM). In most cases, the domain is the same as the AD domain.
NTP Server	Select Use Active Directory Domain to configure an NTP server using the Active Directory DNS. If you need to configure an NTP server using a different address, then you should use the API. See the Cloud Manager API Developer Guide for details.

5. On the Usage Profile, Disk Type, and Tiering Policy page, choose whether you want to enable storage efficiency features, choose a disk type, and edit the S3 tiering policy, if needed.

For help, refer to the following:

- [Understanding volume usage profiles](#)
- [Sizing your system in AWS](#)
- [Sizing your system in Azure](#)
- [Data tiering overview](#)

6. Click **Go**.

Result

Cloud Volumes ONTAP provisions the volume.

After you finish

If you provisioned a CIFS share, give users or groups permissions to the files and folders and verify that those users can access the share and create a file.

If you want to apply quotas to volumes, you must use System Manager or the CLI. Quotas enable you to restrict or track the disk space and number of files used by a user, group, or qtree.

Provisioning volumes on the second node in an HA configuration

By default, Cloud Manager creates volumes on the first node in an HA configuration. If you need an active-active configuration, in which both nodes serve data to clients, you must create aggregates and volumes on the second node.

Steps

1. On the Working Environments page, double-click the name of the Cloud Volumes ONTAP working environment on which you want to manage aggregates.
2. Click the menu icon and then click **Advanced > Advanced allocation**.
3. Click **Add Aggregate** and then create the aggregate.
4. For Home Node, choose the second node in the HA pair.
5. After Cloud Manager creates the aggregate, select it and then click **Create volume**.
6. Enter details for the new volume, and then click **Create**.

After you finish

You can create additional volumes on this aggregate if required.



For HA pairs deployed in multiple AWS Availability Zones, you must mount the volume to clients by using the floating IP address of the node on which the volume resides.

Creating aggregates

You can create aggregates yourself or let Cloud Manager do it for you when it creates volumes. The benefit of creating aggregates yourself is that you can choose the underlying disk size, which enables you to size your aggregate for the capacity or the performance that you need.

Steps

1. On the Working Environments page, double-click the name of the Cloud Volumes ONTAP instance on which you want to manage aggregates.
2. Click the menu icon, and then click **Advanced > Advanced allocation**.
3. Click **Add Aggregate** and then specify details for the aggregate.

For help with disk type and disk size, see [Planning your configuration](#).

4. Click **Go**, and then click **Approve and Purchase**.

Provisioning iSCSI LUNs

If you want to create iSCSI LUNs, you need to do so from System Manager.

Before you begin

- The Host Utilities must be installed and set up on the hosts that will connect to the LUN.
- You must have recorded the iSCSI initiator name from the host. You need to supply this name when you create an igroup for the LUN.
- Before you create volumes in System Manager, you must ensure that you have an aggregate with sufficient space. You need to create aggregates in Cloud Manager. For details, see [Creating aggregates](#).

About this task

These steps describe how to use System Manager for version 9.3 and later.

Steps


1. [Log in to System Manager](#).
2. Click **Storage > LUNs**.
3. Click **Create** and follow the prompts to create the LUN.
4. Connect to the LUN from your hosts.

For instructions, see the [Host Utilities documentation](#) for your operating system.

Tiering inactive data to low-cost object storage

You can reduce storage costs in AWS and Azure by combining an SSD or HDD performance tier for hot data with an object storage capacity tier for inactive data. For a high-level overview, see [Data tiering overview](#).

To set up data tiering, you simply need to do the following:



What's not required for data tiering



- You do not need to install a feature license to enable data tiering.
- You do not need to create the capacity tier (either an S3 bucket or an Azure Blob container). Cloud Manager does that for you.

Configurations that support data tiering

You can enable data tiering when using specific configurations and features:

- Data tiering is supported with Cloud Volumes ONTAP Standard, Premium, and BYOL, starting with version 9.2 in AWS and version 9.4 in Microsoft Azure.
 - Data tiering is not supported with HA pairs in Microsoft Azure.
 - Data tiering is not supported in Azure with the DS3_v2 virtual machine type.
- In AWS, the performance tier can be General Purpose SSDs, Provisioned IOPS SSDs, or Throughput Optimized HDDs.
- In Azure, the performance tier can be Premium SSD managed disks, Standard SSD managed disks, or Standard HDD managed disks.
- Data tiering is supported with encryption technologies.
- Thin provisioning must be enabled on volumes.

Requirements for tiering data in AWS

You must ensure that Cloud Volumes ONTAP has a connection to S3. The best way to provide that connection is by creating a VPC Endpoint to the S3 service. For instructions, see [AWS Documentation: Creating a Gateway Endpoint](#).

When you create the VPC Endpoint, be sure to select the region, VPC, and route table that corresponds to the Cloud Volumes ONTAP instance. You must also modify the security group to add an outbound HTTPS rule that enables traffic to the S3 endpoint. Otherwise, Cloud Volumes ONTAP cannot connect to the S3 service.

If you experience any issues, see [AWS Support Knowledge Center: Why can't I connect to an S3 bucket using a gateway VPC endpoint?](#)

Requirements for tiering data in Microsoft Azure

You do not need to set up a connection between the performance tier and the capacity tier as long as Cloud Manager has the required permissions. Cloud Manager enables a VNet service endpoint for you if the Cloud Manager policy has the appropriate permission:

```
"Microsoft.Network/virtualNetworks/subnets/write",  
"Microsoft.Network/routeTables/join/action",
```

These permissions are included in the latest [Cloud Manager policy](#).

Tiering data on read-write volumes

Cloud Volumes ONTAP can tier inactive data on read-write volumes to cost-effective object storage, freeing up the performance tier for hot data.

Steps

1. In the working environment, create a new volume or change the tier of an existing volume:

Task	Action
Create a new volume	Click Add New Volume .
Modify an existing volume	Select the volume and click Change Disk Type & Tiering Policy .

2. Select the Snapshot Only policy or the Auto policy.

For a description of these policies, see [Data tiering overview](#).

Example



Tiering data to object storage

Volume Tiering Policy

- Auto** - Tiers cold Snapshot copies and cold user data from the active file system to object storage.
- Snapshot Only** - Tiers cold Snapshot copies to object storage
- None** - Data tiering is disabled.

Cloud Manager creates a new aggregate for the volume if a data tiering-enabled aggregate does not already exist.



If you prefer to create aggregates yourself, you can enable data tiering on aggregates when you create them.


Tiering data on data protection volumes

Cloud Volumes ONTAP can tier data from a data protection volume to a capacity tier. If you activate the destination volume, the data gradually moves to the performance tier as it is read.

Steps

1. On the Working Environments page, select the working environment that contains the source volume, and then drag it to the working environment to which you want to replicate the volume.
2. Follow the prompts until you reach the tiering page and enable data tiering to object storage.

Example



For help with replicating data, see [Replicating data to and from the cloud](#).

Changing the tiering level

When you enable data tiering, Cloud Volumes ONTAP tiers inactive data to the S3 *Standard* storage class in AWS or to the *hot* storage tier in Azure. After you deploy Cloud Volumes ONTAP, you can reduce your storage costs by changing the tiering level for inactive data that has not been accessed for 30 days. The access costs are higher if you do access the data, so you must take that into consideration before you change the tiering level.

About this task

The tiering level is system wide—it is not per volume.

In AWS, you can change the tiering level so inactive data moves to one of the following storage classes after 30 days of inactivity:

- Intelligent Tiering
- Standard-Infrequent Access
- One Zone-Infrequent Access

In Azure, you can change the tiering level so inactive data moves to the *cool* storage tier after 30 days of inactivity.

For more information about how tiering levels work, see [Data tiering overview](#).

Steps

1. From the working environment, click the menu icon and then click **Tiering Level**.
2. Choose the tiering level and then click **Save**.

Using Cloud Volumes ONTAP as persistent storage for Kubernetes

Cloud Manager can automate the deployment of [NetApp Trident](#) on Kubernetes clusters so you can use Cloud Volumes ONTAP as persistent storage for containers. Getting started includes a few steps.

If you deploy Kubernetes clusters using the [NetApp Kubernetes Service](#), Cloud Manager can automatically discover the clusters from your NetApp Cloud Central account. If that's the case, skip the first two steps and start with step 3.

1 Verify network connectivity

- A network connection must be available between Cloud Manager and the Kubernetes clusters, and from the Kubernetes clusters to Cloud Volumes ONTAP systems.
- Cloud Manager needs an outbound internet connection to access the following endpoints when installing Trident:

<https://packages.cloud.google.com/yum>
<https://github.com/NetApp/trident/releases/download/>

Cloud Manager installs Trident on a Kubernetes cluster when you connect a working environment to the cluster.

2 Upload Kubernetes configuration files to Cloud Manager

For each Kubernetes cluster, the Cloud Manager Admin needs to upload a configuration file (kubeconfig) that is in YAML format. After you upload the file, Cloud Manager verifies connectivity to the cluster and saves an encrypted copy of the kubeconfig file.

Click **Kubernetes Clusters > Discover > Upload File** and select the kubeconfig file.

The screenshot shows the Cloud Manager interface. On the left, a navigation bar has 'Kubernetes Clusters' highlighted with a red box. Below it, a card with a Kubernetes logo and the text 'Let's discover your first Kubernetes cluster' has a 'Discover' button highlighted with a red box. On the right, the 'Upload Kubernetes Configuration File' page is shown, with an 'Upload File' button highlighted with a red box.

3 Connect your working environments to Kubernetes clusters

From the working environment, click the Kubernetes icon and follow the prompts. You can connect different clusters to different Cloud Volumes ONTAP systems and multiple clusters to the same Cloud Volumes ONTAP system.

You have the option to set the NetApp storage class as the default storage class for the Kubernetes cluster. When a user creates a persistent volume, the Kubernetes cluster can use connected Cloud Volumes ONTAP

systems as the backend storage by default.



4

Start provisioning Persistent Volumes

Request and manage Persistent Volumes using native Kubernetes interfaces and constructs. Cloud Manager creates two Kubernetes storage classes that you can use when provisioning Persistent Volumes:

- **netapp-file**: for binding Persistent Volumes to single-node Cloud Volumes ONTAP systems
- **netapp-file-redundant**: for binding Persistent Volumes to Cloud Volumes ONTAP HA pairs

Cloud Manager configures Trident to use the following provisioning options by default:

- Thin volumes
- The default Snapshot policy
- Accessible Snapshot directory

[Learn more about provisioning your first volume with Trident for Kubernetes](#)

What are the trident_trident volumes?

Cloud Manager creates a volume on the first Cloud Volumes ONTAP system that you connect to a Kubernetes cluster. The name of the volume is appended with "_trident_trident." Cloud Volumes ONTAP systems use this volume to connect to the Kubernetes cluster. You should not delete these volumes.

What happens when you disconnect or remove a Kubernetes cluster?

Cloud Manager enables you to disconnect individual Cloud Volumes ONTAP systems from a Kubernetes cluster. When you disconnect a system, you can no longer use that Cloud Volumes ONTAP system as persistent storage for containers. Existing Persistent Volumes are not deleted.

After you disconnect all systems from a Kubernetes cluster, you can also remove the entire Kubernetes configuration from Cloud Manager. Cloud Manager does not uninstall Trident when you remove the cluster and it does not delete any Persistent Volumes.

Both of these actions are available through APIs only. We plan to add the actions to the interface in a future release.

[Click here for details about the APIs.](#)

Encrypting volumes with NetApp Volume Encryption

NetApp Volume Encryption (NVE) is a software-based technology for encrypting data at rest one volume at a time. Data, Snapshot copies, and metadata are encrypted. Access to the data is given by a unique XTS-AES-256 key, one per volume.

About this task

At this time, Cloud Volumes ONTAP supports NetApp Volume Encryption with an external key management server. An Onboard Key Manager is not supported.

You need to set up NetApp Volume Encryption from the ONTAP CLI. You can then use either the CLI or System Manager to enable encryption on specific volumes. Cloud Manager does not support NetApp Volume Encryption from its user interface and from its APIs.

[Learn more about supported encryption technologies.](#)

Steps

1. Review the list of supported key managers in the [NetApp Interoperability Matrix Tool](#).



Search for the **Key Managers** solution.

2. [Connect to the Cloud Volumes ONTAP CLI.](#)
3. Install a NetApp Volume Encryption license on the Cloud Volumes ONTAP system.

[ONTAP 9 NetApp Encryption Power Guide: Installing the license](#)

4. Install SSL certificates and connect to the external key management servers.

[ONTAP 9 NetApp Encryption Power Guide: Configuring external key management](#)

5. Create a new encrypted volume or convert an existing unencrypted volume using either the CLI or System Manager.

◦ CLI:

- For new volumes, use the **volume create** command with the **-encrypt** parameter.

[ONTAP 9 NetApp Encryption Power Guide: Enabling encryption on a new volume](#)

- For existing volumes, use the **volume encryption conversion start** command.

[ONTAP 9 NetApp Encryption Power Guide: Enabling encryption on an existing volume with the volume encryption conversion start command](#)

- System Manager:

- For new volumes, click **Storage > Volumes > Create > Create FlexVol** and then select **Encrypted**.

[ONTAP 9 Cluster Management using System Manager: Creating FlexVol volumes](#)

- For existing volumes, select the volume, click **Edit**, and then select **Encrypted**.

[ONTAP 9 Cluster Management using System Manager: Editing volume properties](#)

Managing existing storage

Cloud Manager enables you to manage volumes, aggregates, and CIFS servers. It also prompts you to move volumes to avoid capacity issues.




Managing existing volumes

You can manage existing volumes as your storage needs change. You can view, edit, clone, restore, and delete volumes.

Steps

1. On the Working Environments page, double-click the Cloud Volumes ONTAP working environment on which you want to manage volumes.
2. Manage your volumes:

Task	Action
View information about a volume	Select a volume, and then click Info .
Edit a volume (read-write volumes only)	<ol style="list-style-type: none"> a. Select a volume, and then click Edit. b. Modify the volume's Snapshot policy, NFS access control list, or share permissions, and then click Update.
Clone a volume	<ol style="list-style-type: none"> a. Select a volume, and then click Clone. b. Modify the clone name as needed, and then click Clone. <p>This process creates a FlexClone volume. A FlexClone volume is a writable, point-in-time copy that is space-efficient because it uses a small amount of space for metadata, and then only consumes additional space as data is changed or added.</p> <p>To learn more about FlexClone volumes, see the ONTAP 9 Logical Storage Management Guide.</p>

Task	Action
Restore data from a Snapshot copy to a new volume	<ol style="list-style-type: none"> Select a volume, and then click Restore from Snapshot copy. Select a Snapshot copy, enter a name for the new volume, and then click Restore.
Create a Snapshot copy on demand	<ol style="list-style-type: none"> Select a volume, and then click Create a Snapshot copy. Change the name, if needed, and then click Create.
Get the NFS mount command	<ol style="list-style-type: none"> Select a volume, and then click Mount Command. Click Copy.
Change the underlying disk type	<ol style="list-style-type: none"> Select a volume, and then click Change Disk Type & Tiering Policy. Select the disk type, and then click Change. <p> Cloud Manager moves the volume to an existing aggregate that uses the selected disk type or it creates a new aggregate for the volume.</p>
Change the tiering policy	<ol style="list-style-type: none"> Select a volume, and then click Change Disk Type & Tiering Policy. Click Edit Policy. Select a different policy and click Change. <p> Cloud Manager moves the volume to an existing aggregate that uses the selected disk type with tiering, or it creates a new aggregate for the volume.</p>
Enable or disable sync to S3 for a volume	<p>Select a volume and then click Sync to S3 or Delete Sync Relationship.</p> <p> The sync to S3 feature must be enabled before you can use these options. For instructions, see Syncing data to AWS S3</p>
Delete a volume	<ol style="list-style-type: none"> Select a volume, and then click Delete. Click Delete again to confirm.

Managing existing aggregates

Manage aggregates yourself by adding disks, viewing information about the aggregates, and by deleting them.

Before you begin

If you want to delete an aggregate, you must have first deleted the volumes in the aggregate.


About this task

If an aggregate is running out of space, you can move volumes to another aggregate by using OnCommand

System Manager.

Steps

1. On the Working Environments page, double-click the Cloud Volumes ONTAP working environment on which you want to manage aggregates.
2. Click the menu icon and then click **Advanced > Advanced allocation**.
3. Manage your aggregates:

Task	Action
View information about an aggregate	Select an aggregate and click Info .
Create a volume on a specific aggregate	Select an aggregate and click Create volume .
Add disks to an aggregate	<ol style="list-style-type: none">a. Select an aggregate and click Add AWS disks or Add Azure disks.b. Select the number of disks that you want to add and click Add. <p> All disks in an aggregate must be the same size.</p>
Delete an aggregate	<ol style="list-style-type: none">a. Select an aggregate that does not contain any volumes and click Delete.b. Click Delete again to confirm.

Modifying the CIFS server

If you change your DNS servers or Active Directory domain, you need to modify the CIFS server in Cloud Volumes ONTAP so that it can continue to serve storage to clients.

Steps

1. From the working environment, click the menu icon and then click **Advanced > CIFS setup**.
2. Specify settings for the CIFS server:

Task	Action
DNS Primary and Secondary IP Address	The IP addresses of the DNS servers that provide name resolution for the CIFS server. The listed DNS servers must contain the service location records (SRV) needed to locate the Active Directory LDAP servers and domain controllers for the domain that the CIFS server will join.
Active Directory Domain to join	The FQDN of the Active Directory (AD) domain that you want the CIFS server to join.
Credentials authorized to join the domain	The name and password of a Windows account with sufficient privileges to add computers to the specified Organizational Unit (OU) within the AD domain.
CIFS server NetBIOS name	A CIFS server name that is unique in the AD domain.

Task	Action
Organizational Unit	The organizational unit within the AD domain to associate with the CIFS server. The default is CN=Computers.
DNS Domain	The DNS domain for the Cloud Volumes ONTAP storage virtual machine (SVM). In most cases, the domain is the same as the AD domain.
NTP Server	Select Use Active Directory Domain to configure an NTP server using the Active Directory DNS. If you need to configure an NTP server using a different address, then you should use the API. See the Cloud Manager API Developer Guide for details.

3. Click **Save**.

Result

Cloud Volumes ONTAP updates the CIFS server with the changes.

Moving a volume to avoid capacity issues

Cloud Manager might display an Action Required message that says moving a volume is necessary to avoid capacity issues, but that it cannot provide recommendations to correct the issue. If this happens, you need to identify how to correct the issue and then move one or more volumes.

Steps

1. [Identify how to correct the issue](#).
2. Based on your analysis, move volumes to avoid capacity issues:
 - [Move volumes to another system](#).
 - [Move volumes to another aggregate on the same system](#).


Identifying how to correct capacity issues

If Cloud Manager cannot provide recommendations for moving a volume to avoid capacity issues, you must identify the volumes that you need to move and whether you should move them to another aggregate on the same system or to another system.

Steps

1. View the advanced information in the Action Required message to identify the aggregate that has reached its capacity limit.

For example, the advanced information should say something similar to the following: Aggregate aggr1 has reached its capacity limit.
2. Identify one or more volumes to move out of the aggregate:
 - a. In the working environment, click the menu icon, and then click **Advanced > Advanced allocation**.
 - b. Select the aggregate, and then click **Info**.
 - c. Expand the list of volumes.



d. Review the size of each volume and choose one or more volumes to move out of the aggregate.

You should choose volumes that are large enough to free space in the aggregate so that you avoid additional capacity issues in the future.

3. If the system has not reached the disk limit, you should move the volumes to an existing aggregate or a new aggregate on the same system.

For details, see [Moving volumes to another aggregate to avoid capacity issues](#).

4. If the system has reached the disk limit, do any of the following:

- a. Delete any unused volumes.
- b. Rearrange volumes to free space on an aggregate.

For details, see [Moving volumes to another aggregate to avoid capacity issues](#).

c. Move two or more volumes to another system that has space.

For details, see [Moving volumes to another system to avoid capacity issues](#).

Moving volumes to another system to avoid capacity issues

You can move one or more volumes to another Cloud Volumes ONTAP system to avoid capacity issues. You might need to do this if the system reached its disk limit.

About this task

You can follow the steps in this task to correct the following Action Required message:

Moving a volume is necessary to avoid capacity issues; however, Cloud Manager cannot perform this action for you because the system has reached the disk limit.

Steps

1. Identify a Cloud Volumes ONTAP system that has available capacity, or deploy a new system.
2. Drag and drop the source working environment on the target working environment to perform a one-time data replication of the volume.

For details, see [Replicating data between systems](#).

3. Go to the Replication Status page, and then break the SnapMirror relationship to convert the replicated volume from a data protection volume to a read/write volume.

For details, see [Managing data replication schedules and relationships](#).

4. Configure the volume for data access.

For information about configuring a destination volume for data access, see the [ONTAP 9 Volume Disaster Recovery Express Guide](#).

5. Delete the original volume.

For details, see [Managing existing volumes](#).

Moving volumes to another aggregate to avoid capacity issues

You can move one or more volumes to another aggregate to avoid capacity issues.

About this task

You can follow the steps in this task to correct the following Action Required message:

```
Moving two or more volumes is necessary to avoid capacity issues; however,
Cloud Manager cannot perform this action for you.
```

Steps

1. Verify whether an existing aggregate has available capacity for the volumes that you need to move:
 - a. In the working environment, click the menu icon, and then click **Advanced > Advanced allocation**.
 - b. Select each aggregate, click **Info**, and then view the available capacity (aggregate capacity minus used aggregate capacity).

aggr1

Aggregate Capacity: 442.94 GB

Used Aggregate Capacity: 105.66 GB

2. If needed, add disks to an existing aggregate:
 - a. Select the aggregate, and then click **Add disks**.
 - b. Select the number of disks to add, and then click **Add**.
3. If no aggregates have available capacity, create a new aggregate.

For details, see [Creating aggregates](#).

4. Use System Manager or the CLI to move the volumes to the aggregate.
5. In most situations, you can use System Manager to move volumes.

For instructions, see the [ONTAP 9 Volume Move Express Guide](#).

Provisioning NFS volumes from the Volume View

Changing to the Volume View

Cloud Manager provides two management views: the Storage System View for managing storage systems across a hybrid cloud and the Volume View for creating volumes in AWS without having to manage storage systems. You can switch between these views, but those instances should be rare because a single view should meet your needs.

For more information about the Volume View, see [Simplified storage management using the Volume View](#).

Steps

1. In the upper right of the Cloud Manager console, click the menu, and then click **View Selection**.
2. On the View Selection page, select **Storage System View**, and then click **Switch**.

Result

Cloud Manager switches to the Volume View.

Creating and mounting NFS volumes

You can use Cloud Manager to create NFS volumes that provide enterprise-class features on top of AWS storage.

Creating NFS volumes

You can create a volume attached to a single AWS instance or to an instance that is mirrored to another instance to provide high availability.

Steps

1. In the Volumes tab, click **Create New Volume**.
2. On the Create New Volume page, select a volume type:

Option	Description
Create Volume	Creates a volume attached to a single AWS instance.
Create HA volume	Creates a volume attached to a single AWS instance and mirrored to another instance to provide high availability in case of failures. Click the Info icon to see additional details about the instances required for an HA volume.

3. If you chose Create Volume, specify details for your first volume, and then click **Create**.

The following table describes fields for which you might need guidance:

Field	Description
Size	The maximum size for the volume depends on the capacity available in existing storage systems. Thin provisioning is automatically enabled on the volume, which enables you to create a volume that is bigger than the physical storage currently available to it. Instead of preallocating storage space, space is allocated to each volume as data is written.
AWS Disk Type	You should choose the disk that meets your requirements for both performance and cost. <ul style="list-style-type: none"> • General Purpose SSD disks balance cost and performance for a broad range of workloads. Performance is defined in terms of IOPS. • Throughput Optimized HDD disks are for frequently accessed workloads that require fast and consistent throughput at a lower price. • Cold HDD disks are meant for backups, or infrequently accessed data, because the performance is very low. Like Throughput Optimized HDD disks, performance is defined in terms of throughput. <p>For more details, refer to AWS Documentation: EBS Volume Types.</p>

The following image shows the Create Volume page filled out:

The screenshot shows the AWS Create Volume page with the following details:

Details	Location
<p>Volume Name: <input type="text" value="vol1"/></p> <p>Size (GB): <input type="text" value="500"/></p> <p>AWS Disk Type: <input type="text" value="General Purpose (SSD)"/></p>	<p>AWS Region: US East N. Virginia</p> <p>VPC: vpc-a6c1eac2 172.32.0.0/16</p> <p>Subnet: 172.32.0.0/24</p>

4. If you chose Create HA volume, specify details for the volume, and then click **Create**.

The following table describes fields for which you might need guidance:

Field	Description
Size	The maximum size for the volume depends on the capacity available in existing storage systems. Thin provisioning is automatically enabled on the volume, which enables you to create a volume that is bigger than the physical storage currently available to it. Instead of preallocating storage space, space is allocated to each volume as data is written.

Field	Description
AWS Disk Type	<p>You should choose the disk that meets your requirements for both performance and cost.</p> <ul style="list-style-type: none"> • General Purpose SSD disks balance cost and performance for a broad range of workloads. Performance is defined in terms of IOPS. • Throughput Optimized HDD disks are for frequently accessed workloads that require fast and consistent throughput. <p>For more details, refer to AWS Documentation: EBS Volume Types.</p>
Location	You should choose a VPC that includes three subnets in three separate Availability Zones.
Nodes and Mediator	If possible, Cloud Manager chooses separate Availability Zones for each instance because it is the supported and optimal configuration.
Floating IP	The IP addresses must be outside of the CIDR block for all VPCs in the region.
Route Table	<p>If you have more than one route table, it is very important to select the correct route tables. Otherwise, some clients might not have access to the HA pair.</p> <p>For more details, refer to AWS Documentation: Route Tables.</p>

The following image shows the Nodes and Mediator page. Each instance is in a separate Availability Zone.

Nodes & Mediator Edit			
Node 1	Availability Zone us-east-1d	Subnet 172.31.0.0/20	
Node 2	Availability Zone us-east-1c	Subnet 172.31.16.0/20	
Mediator	Availability Zone us-east-1b	Subnet 172.31.32.0/20	Key Pair EranVirginia

Result

Cloud Manager creates the volume on an existing system or on a new system. If a new system is required, creating the volume can take approximately 25 minutes.

Mounting volumes to Linux hosts

After you create a volume, you should mount it to your hosts so that they can access the volume.

Steps

1. In the Volumes tab, place your mouse cursor over the volume, select the menu icon, and then click **Mount**.
2. Click **Copy**.
3. On your Linux hosts, modify the copied text by changing the destination directory, and then enter the command to mount the volume.

Managing NFS volumes

You can manage NFS volumes by cloning them, managing data access, changing the underlying disk type, and more.

Cloning volumes

If you need an instantaneous copy of your data without using a lot of disk space, you can create a clone of an existing volume.

About this task

The cloned volume is a writable, point-in-time copy that is space-efficient because it uses a small amount of space for metadata, and then only consumes additional space as data is changed or added.

Steps

1. In the Volumes tab, place your mouse cursor over the volume, select the menu icon, and then click **Clone**.
2. Modify the name of the cloned volume, if needed, and then click **Clone**.

Result

Cloud Manager creates a new volume that is a clone of an existing volume.

Managing data access to volumes

When you create a volume, Cloud Manager makes the volume available to all EC2 instances in the VPC in which the volume was created. You can modify this default value if you need to restrict data access to the volume.

Steps

1. In the Volumes tab, place your mouse cursor over the volume, select the menu icon, and then click **Manage Access**.
2. Modify the volume access list, and then click **Save**.

Changing the underlying AWS disk for a volume

You can change the underlying AWS disk that a volume uses to provide storage. For example, if higher performance is needed, you can change from a Throughput Optimized HDD to a General Purpose SSD.

Steps

1. In the Volumes tab, place your mouse cursor over the volume, select the menu icon, and then click **Change Disk**.
2. Select the AWS disk type and click **Change**.

Result

Cloud Manager moves the volume to an existing aggregate that uses the selected disk type or it creates a new aggregate for the volume.

Viewing and modifying AWS resources


When you create a new volume, Cloud Manager allocates the AWS instances and EBS storage required for that volume. If required, you can view details about AWS instances and EBS storage, change instance types, and turn instances off and on.

Steps

1. Click **AWS Resources**.

The list of AWS instances displays. You can view details such as instance type, AWS location, and the volumes attached to the instance.

2. If required, select the menu icon next to the Status column, and then choose one of the available actions:



Deleting volumes

You can delete volumes that you no longer need.

Steps

1. In the Volumes tab, place your mouse cursor over the volume, select the menu icon, and then click **Delete**.
2. Click **Delete** to confirm that you want to delete the volume.

Copyright information

Copyright © 2024 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.