



Manage volumes

Workload Factory for NetApp Console

NetApp
October 07, 2026

This PDF was generated from <https://docs.netapp.com/us-en/data-services-workloads/fsx-storage/manage-volume-capacity.html> on October 07, 2026. Always check docs.netapp.com for the latest.

Table of Contents

Manage volumes	1
Manage volume capacity in NetApp Workload Factory	1
Edit volume autogrow	1
Adjust volume capacity manually	1
Rebalance volume capacity for FlexVol and FlexGroup volumes	2
Set volume latency thresholds in NetApp Workload Factory	5
Manage volume settings in NetApp Workload Factory	6
Manage volume tags	6
Update volume storage efficiency	7
Manage immutable files for a volume in NetApp Workload Factory	7
Related information	8
Manage FSx for ONTAP cache volumes with NetApp Workload Factory	8
Edit the cache volume name	8
Adjust the capacity of a cache volume	9
Edit the cache volume export policy	9
Change the caching method for a cache volume	9
Prepopulate a cache volume	10
Delete a cache volume	10
Change the tiering policy of a volume in NetApp Workload Factory	10
Manage the NFS export policy for a volume in NetApp Workload Factory	12
Manage the SMB/CIFS share for a volume in NetApp Workload Factory	13
Create an SMB/CIFS share for a volume	13
Change an SMB/CIFS share for a volume	14
Delete an SMB/CIFS share for a volume	14
Manage S3 access points	14
Create S3 access points for a volume in NetApp Workload Factory	14
Manage the S3 access points for a volume in NetApp Workload Factory	16
Set up the Journal table infrastructure for NetApp Workload Factory	17
Split a cloned volume in NetApp Workload Factory	26
Delete a volume in NetApp Workload Factory	27

Manage volumes

Manage volume capacity in NetApp Workload Factory

Volume capacity should be adjusted as needs change. You can adjust volume size manually or turn on volume autogrow to let Workload Factory manage volume capacity for you. Check the balance of FlexVol or FlexGroup volume capacity and rebalance volume capacity to spread files evenly across all FlexVol volumes in a node or across all constituents so that all nodes participate in the workload of a single FlexGroup volume.

Edit volume autogrow

Turn on volume autogrow to let Workload Factory manage volume capacity for you. You can disable it at any time.



Volume autogrow isn't supported for iSCSI volumes.

Before you begin

To enable volume autogrow, you must associate a link. [Learn how to associate an existing link or to create and associate a new link](#). After the link associates, return to this operation.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Volumes explorer**.
3. Select the actions menu for the volume you want to modify.
4. Select **Basic actions**, then **Set volume autogrow**.
5. In the Set autogrow dialog, enable volume autogrow to automatically expand the volume capacity until the volume reaches the maximum size. This feature accommodates increasing data usage, ensuring uninterrupted operations.

Specify the maximum volume growth size and unit. The maximum growth size cannot be smaller than the current volume size.

6. Select **Apply**.

Adjust volume capacity manually

You can adjust volume capacity by increasing or decreasing the provisioned size of a volume. The following table shows the minimum and maximum volume sizes by volume type:

Volume type	Minimum size	Maximum size
FlexVol volume	20 MiB	300 TiB
FlexGroup volume	800 GiB	2 PiB

For an iSCSI LUN, increasing the size of the volume also increases the size of the host LUN. After you increase volume capacity, follow the procedure provided by your host operating system to discover the new

size of the LUN and expand the file system of the LUN.

Decreasing volume size is supported only for NFS and SMB/CIFS volumes.

Before you begin

To adjust volume capacity, you must associate a link. [Learn how to associate an existing link or to create and associate a new link](#). After the link associates, return to this operation.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Volumes explorer**.
3. Select the actions menu for the volume you want to modify.
4. Select **Basic actions**, then **Adjust volume capacity**.
5. In the Adjust volume capacity dialog, set the **Provisioned capacity** and unit.
6. Select **Adjust** to apply the changes.

Rebalance volume capacity for FlexVol and FlexGroup volumes

Rebalancing volume capacity is supported for FlexVol volumes and FlexGroup volumes. Rebalancing a volume redistributes the capacity when imbalances develop over time due to the addition of new files and file growth. After you manually start the rebalance operation, we select the files and move them automatically and non-disruptively. Volume transfer operations consume file system resources.

Each volume type and rebalancing operations differ as follows.

FlexVol volumes

FlexVol volumes are logical containers that offer flexibility in managing data, allowing for expansion, contraction, movement, and efficient copying. They can be used with NAS and SAN environments.

A FlexVol volume can be balanced in relation to other FlexVol volumes within one node in an FSx for ONTAP file system. If the file system has only a single FlexVol volume, then rebalancing isn't possible. When the file system has more than one FlexVol volume per node and a single FlexVol volume is selected, the FlexVol volume is balanced in the context of all FlexVols but only the selected volume is allowed to move.

FlexGroup volumes

FlexGroup volumes, on the other hand, are scalable NAS containers designed for high performance and automatic load distribution. They consist of multiple member volumes (constituents) that share traffic transparently. FlexGroup volumes provide massive capacity, exceeding FlexVol limits, with up to 60PB capacity and 400 billion files. They simplify management by offering a single namespace container.

Capacity is spread across a number of constituents in a scale-out FSx for ONTAP file system with two or more high availability (HA) pairs. Each constituent is a container that dictates the maximum single file size. FSx for ONTAP spreads files across all constituents in an even way so all nodes participate in the workload of a single FlexGroup volume.

When the constituents aren't distributed evenly across all nodes, FlexGroup volume performance decreases.

Checking the balance of FlexGroup volume capacity includes assessing the current layout of constituents. When you rebalance the volume's capacity, NetApp Workload Factory designs a new constituent layout with an even number of constituents to spread the data evenly across all HA pairs. The service executes the rebalance plan which in turn improves read and write operations.



Rebalancing isn't supported for SAN volumes like iSCSI and NVMe.

Check the balance of your volumes

Check the balance of FlexVol or FlexGroup volumes in an FSx for ONTAP file system.

Before you begin

- FlexGroup volume balance is available only for FSx for ONTAP file systems using a scale-out deployment with at least two HA pairs.
- To check the balance of a volume, you must [create and associate a link or associate an existing link to the file system](#). After the link associates, return to this operation.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads** > **Storage** > **FSx for ONTAP**.
3. From **FSx for ONTAP**, select the actions menu of the file system that contains volumes to rebalance and then select **Manage**.
4. From the file system overview, select the **Volumes** tab.
5. From the Volumes tab, select **Check volume balance** at the top of the table.
6. In the Volume balance window, review the balance status of:
 - FlexGroup volumes
 - FlexVol volumes

When a volume is unbalanced, consider [rebalancing it](#).

Rebalance volume capacity

Rebalance one or more unbalanced volumes.



A Workload Factory admin can [stop rebalancing](#) during the operation.

Before you begin

- [Check the balance of a volume](#) before rebalancing volumes.
- To rebalance a volume, you must [create and associate a link or associate an existing link to the file system](#). After the link associates, return to this operation.
- Note that existing snapshots on volumes you rebalance become partial and cannot be used to restore volume data, but new snapshots taken after rebalancing can be used to restore volume data.
- FlexVol volumes are best rebalanced altogether to balance all volume resources evenly. Deselected volumes don't actively participate in the balancing procedure.

FlexVol volume

A FlexVol volume can be balanced in relation to other FlexVol volumes within one node in an FSx for ONTAP file system. When the file system has more than one FlexVol volume per node and a single FlexVol volume is selected, then the FlexVol volume is balanced in the context of all FlexVols but only the selected volume is allowed to move.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From **FSx for ONTAP**, select the actions menu of the file system that contains the volume to rebalance and then select **Manage**.
4. From the file system overview, select the **Volumes** tab.
5. From the Volumes tab, select **Check volume balance** at the top of the table.
6. In the Volumes balance window, optionally select **Data distribution** in the FlexVol balance summary to view used capacity per aggregate.
7. Select **Rebalance** to rebalance one or more unbalanced volumes.
8. In the Rebalance wizard, follow the steps.
 - a. **Max transfer rate**: Optional. Disabled by default. Enable throttling to limit the bandwidth of a volume move on your file system and to slow outgoing volume replication traffic.

Enter the throttle value in MB/s.

Select **Next**.
 - b. Review the current and proposed layouts of all FlexVol volumes, and then select **Next**.
 - c. Carefully review what will happen and the note before beginning the rebalance operation.
9. Select **Rebalance**.

Result

The FlexVol volume is rebalanced. When the operation completes, the file system will be throttled back to the original value.

FlexGroup volume

Data redistributes across member volumes to rebalance the FlexGroup volume. Based on your chosen layout, the rebalance operation might add FlexGroup member volumes and increase the size of provisioned volumes.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From **FSx for ONTAP**, select the actions menu of the file system that contains the volume to rebalance and then select **Manage**.
4. From the file system overview, select the **Volumes** tab.
5. From the Volumes tab, select **Check FlexGroup balance** at the top of the table.
6. In the FlexGroup balance window, select **Rebalance** to rebalance one or more unbalanced volumes.
7. In the Rebalance wizard, select the data distribution layout that you prefer.
 - **Performance-optimized** (recommended): increases the number of FlexGroup member volumes and the provisioned size of the volume. Follows NetApp best practice.
 - **Restricted**: supports volumes in a replication relationship. The number of FlexGroup member volumes and the size of provisioned volumes remains the same. Selected by default if all selected

volumes participate in a replication relationship.

- **Manual:** Select the desired number of FlexGroup member volumes per HA pair. Depending on your selection, the number of FlexGroup member volumes and the provisioned size of the volume might increase.
8. **Throttling:** Optional. Disabled by default. Enable throttling to limit the bandwidth of a volume move on your file system and to slow outgoing volume replication traffic.

Enter the throttle value in MB/s.

9. Select a layout comparison view and then select **Next**.
- Volume layout comparison
 - FSx for ONTAP layout comparison
10. Optionally, download a list of volume moves before rebalancing.
11. Select **Rebalance**.

Result

FlexGroup member volumes are moved one at a time during rebalancing. When the operation completes, the file system will be throttled back to the original value.

Stop a volume rebalance operation

Stop a rebalance operation at any time; it isn't disruptive. Stopping the operation aborts active volume moves.

You can start another rebalance operation later.

Steps

1. After you begin the rebalance operation, from the Volume balance page, select **Stop rebalancing**.
2. In the Stop rebalancing dialog, select **Stop**.

Result

The volume rebalance operation stops and active volume moves abort.

Set volume latency thresholds in NetApp Workload Factory

Set latency thresholds for volumes to receive alerts when read or write latency exceeds specified values.

About this task

You can configure latency thresholds on a per-volume basis to monitor performance and receive alerts when latency values exceed your defined limits. Workload Factory monitors both read and write latency and generates alerts when either threshold is exceeded.

When you set latency thresholds, Workload Factory monitors the following:

- **Read latency:** Triggers an alert when read operations exceed the specified threshold in milliseconds.
- **Write latency:** Triggers an alert when write operations exceed the specified threshold in milliseconds.

Latency monitoring uses the following default parameters:

- **IOPS threshold:** 100 IOPS (read and write)
- **Time range:** 10 minutes

An alert is generated when the latency threshold is exceeded and the IOPS threshold is met over the specified time range.



You can view latency events and receive alerts without a link connection. However, to analyze events and receive AI-powered insights and recommendations, you must [connect to the file system using a link](#).

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Volumes explorer**.
3. Select the actions menu of the volume to set latency thresholds for.
4. Select **Alerts actions**, then **Set latency threshold**.
5. In the Set latency threshold dialog, configure the following:
 - a. **Read latency:** Enter the threshold value in milliseconds (ms) for read operations.
 - b. **Write latency:** Enter the threshold value in milliseconds (ms) for write operations.
6. Select **Save** to apply the latency thresholds.

Result

Workload Factory monitors the volume and generates alerts when read or write latency exceeds the configured thresholds. You can view latency events from the [Analysis screen](#).

Related information

- [Analyze FSx for ONTAP events in NetApp Workload Factory](#)
- [Adjust volume capacity in Workload Factory](#)

Manage volume settings in NetApp Workload Factory

Customize FSx for ONTAP volume settings in NetApp Workload Factory to fit your organization's needs. You can manage volume tags to categorize resources, update storage efficiency settings to optimize space and cost savings, and manage immutable files settings to control data retention and protect against unauthorized changes.

Manage volume tags

Tags can help you categorize your resources. You can add, edit, and remove volume tags at any time for FSx for ONTAP volumes in NetApp Workload Factory.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Volumes explorer**.
3. Select the actions menu for the volume to modify tags for.

4. Select **Basic actions** then **Edit volume tags**.
5. On the Edit volume tags page, add, edit, or remove tags.

The maximum number of tags you can apply to a volume is 50.

6. Select **Apply**.

Update volume storage efficiency

Update the storage efficiency setting after volume creation.

The storage efficiency feature includes deduplication, data compression, and data compaction to achieve optimal space savings on a FlexVol volume. Deduplication eliminates duplicate data blocks. Data compression compresses the data blocks to reduce the amount of physical storage that is required. Data compaction stores more data in less space to increase storage efficiency.

If you chose not to enable storage efficiency when you created a volume, you can enable the setting for potential space and cost savings at any time.

Volumes use thin provisioning whether you enable or disable storage efficiency.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Volumes explorer**.
3. Select the actions menu of the volume to update the storage efficiency setting for.
4. Select **Advanced actions**, then **Set storage efficiency**.
5. Choose to enable or disable volume storage efficiency.
6. Select **Apply** to save the change.

Manage immutable files for a volume in NetApp Workload Factory

You can update certain immutable files settings for a volume when the feature is enabled, such as the retention policy and periods, the autocommit period, and volume append mode.

Note that enabling immutable files is only possible during [volume creation](#).

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Volumes explorer**.
3. Select the actions menu of the volume to update.
4. Select **Data protection actions**, then **Manage immutable files**.
5. On the Manage Immutable files page, you can update the following:
 - **Retention period**: select **Unspecified** or **Specify period**.
 - **Unspecified**: The default minimum period is "0" years and the default maximum period is "30 years".
 - **Specify period**: Option to define the retention policy, minimum and maximum periods, the

autocommit feature, and the volume append mode feature. Provide the following details:

- **Retention policy:** This period must be greater than or equal to the minimum retention period and less than or equal to the maximum retention period.
- **Minimum and maximum periods:** Set the minimum and maximum periods to commit files in this volume to an immutable WORM state.
- **Autocommit:** enable or disable the feature to automatically commit files to WORM that haven't been modified during the Autocommit period.
- **Privileged delete:** Enable or disable the feature. Enabling the feature allows a SnapLock administrator to delete an unexpired WORM volume. This feature is only supported in Enterprise retention mode.
- **Volume append mode:** enable or disable the feature. Enabling volume append mode enables you to add new content to WORM files.

6. Click **Apply**.

Related information

- [Create a volume in Workload Factory](#)
- [Manage volume capacity in Workload Factory](#)
- [Change volume tiering policy in Workload Factory](#)
- [Manage S3 access points in Workload Factory](#)

Manage FSx for ONTAP cache volumes with NetApp Workload Factory

Use the NetApp Workload Factory console to manage cache volumes for FSx for ONTAP file systems. Caching, a method for temporarily storing data, improves data access performance by reducing retrieval time. You can edit the cache name, adjust capacity, change the export policy, select a caching method, pre-populate the cache, or delete cache volumes.

About this task

You can manage cache volumes that are associated with cache relationships in the NetApp Workload Factory console.

Before you begin

- You must associate a link to manage cache volumes and relationships. [Learn how to associate an existing link or to create and associate a new link](#). After you associate the link, return to this operation.
- You must have an existing cache volume to edit.

Edit the cache volume name

Change the name of an existing cache volume at any time.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads** > **Storage** > **FSx for ONTAP**.

3. From FSx for ONTAP, select the actions menu of the file system with the cache volume and then select **Manage**.
4. From the file system overview, select the **Cache relationships** tab.
5. Select the actions menu for the cache volume you want to modify and then select **Edit cache name**.
6. In the **Edit cache name** dialog, enter the new name for the cache volume and then select **Apply**.

Adjust the capacity of a cache volume

You can adjust the capacity of an existing cache volume at any time.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From FSx for ONTAP, select the actions menu of the file system with the cache volume and then select **Manage**.
4. From the file system overview, select the **Cache relationships** tab.
5. Select the actions menu for the cache volume you want to modify and then select **Adjust cache capacity**.
6. In the **Adjust cache capacity** dialog, enter the new capacity for the cache volume by percentage or by unit and then select **Apply**.

Edit the cache volume export policy

Change the mount path or the export policy assigned to an existing cache volume.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From FSx for ONTAP, select the actions menu of the file system with the cache volume and then select **Manage**.
4. From the file system overview, select the **Cache relationships** tab.
5. Select the actions menu for the cache volume you want to modify and then select **Edit export policy**.
6. In the **Edit export policy** dialog, change the mount path or select a different export policy to assign to the cache volume.
7. Select **Apply**.

Change the caching method for a cache volume

You can change how the cache works for an existing cache volume to write-around or write-back.

Learn more about [write modes](#).

Steps

1. Log in to the [NetApp Console](#).
- 2.

Select the menu , **Workloads > Storage > FSx for ONTAP**.

3. From FSx for ONTAP, select the actions menu of the file system with the cache volume and then select **Manage**.
4. From the file system overview, select the **Cache relationships** tab.
5. Select the actions menu for the cache volume you want to modify and then select **Change caching method**.
6. In the **Change caching method** dialog, select the new caching method and then select **Apply**.

Prepopulate a cache volume

Fill the cache volume with data from the origin volume before you use it to make cached data available faster.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From FSx for ONTAP, select the actions menu of the file system with the cache volume and then select **Manage**.
4. From the file system overview, select the **Cache relationships** tab.
5. Select the actions menu for the cache volume you want to modify and then select **Prepopulate cache**.
6. In the **Prepopulate cache** dialog, specify the path to the data set to use for prepopulation and then select **Apply**.

Delete a cache volume

When you delete a cache volume, you remove its cache relationship. Cached data is no longer available.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From FSx for ONTAP, select the actions menu of the file system with the cache volume and then select **Manage**.
4. From the file system overview, select the **Cache relationships** tab.
5. Select the actions menu for the cache volume you want to delete and then select **Delete cache volume**.
6. In the **Delete cache volume** dialog, confirm the deletion and then select **Delete**.

Change the tiering policy of a volume in NetApp Workload Factory

In NetApp Workload Factory, you can change the tiering policy to automatically re-allocate data from the high-performance primary storage tier to the secondary capacity pool storage tier.

About this task

You can change the tiering policy of a volume at any time. The tiering policy is defined per volume.

Deciding where your data is stored has implications for your cost savings.

FSx for ONTAP has two tiers for storing volume data:

- **SSD storage tier:** This primary storage tier is for the data you access most frequently, also known as *hot* data. Storing data in the primary storage tier is more expensive than in the secondary storage tier.
- **Capacity pool storage tier:** This secondary storage tier is for archived data or infrequently accessed data, also known as *cold* data.

Refer to [Managing storage capacity](#) in AWS for FSx for NetApp ONTAP documentation for more information about storage tiers.

Before you begin

Review the available tiering policies before you change the tiering policy.

- **Balanced (Auto):** default tiering policy when creating a volume using the user interface. Keeps frequently accessed data in the SSD storage tier and tiers infrequently accessed data and snapshots to the capacity pool storage tier after the cooling period ends. Recommended for general primary workloads.
- **Cost-optimized (All):** Tiers all snapshots and data to the capacity pool storage tier. Recommended for secondary targets.
- **Performance optimized (Snapshots only):** Tiers only snapshot data to the capacity pool storage tier. Recommended for low-latency workloads such as mission-critical databases.
- **None:** Keeps volume data in the SSD storage tier, preventing it from being moved to the capacity pool storage tier.

Note that some tiering policies have an associated minimum cooling period which sets the time, or *cooling days*, that user data in a volume must remain inactive for the data to be considered "cold" and moved to the capacity pool storage tier. The cooling period starts when data is written to the disk.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Volumes explorer > FSx for ONTAP**.
3. Select the actions menu of the volume to change the tiering policy for.
4. Select **Advanced actions**, then **Change tiering policy**.
5. On the Change tiering policy page, select to copy the tiering policy of the source volume or select one of the following tiering policies:
 - **Balanced (Auto):** Enter the number of cooling days.
 - **Cost-optimized (All)**
 - **Performance-optimized (Snapshots only):** Enter the number of cooling days.
 - **None**
6. Select **Apply**.

Manage the NFS export policy for a volume in NetApp Workload Factory

Manage the NFS export policy for a volume that uses NFSv3 or NFSv4.1 protocol types in NetApp Workload Factory.

About this task

Managing a volume's export policy involves adding export policy rules that detail client specifications, access control, super user access, and NFS version. You can add more than one export policy and prioritize them.

Before you begin

Determine the client specifications for the export policy rules. Valid values for the client specification are as follows:

- IP addresses
- IP addresses with subnet masks
- IP addresses with a network mask
- A netgroup name preceded by the "@" character
- A domain name preceded by a period "."
- Host names

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Volumes explorer**.
3. Select the actions menu for the volume to change the NFS export policy for.
4. Select **Advanced actions**, then **Edit NFS export policy**.
5. On the Edit NFS export policy page, provide the following:
 - a. **Access control**: Select **Custom export policy** or **Existing export policy**.

Alternatively, you can select **No access to the volume**.
 - b. **Export policy name**: Optionally, enter a name for the export policy.
 - c. **Add export policy rule**: Provide the following details and rank the policies starting with #1 as the priority rule:
 - i. **Client specification**: Separate multiple values with commas.
 - ii. **Access control**: Select **Read/Write**, **Read only**, or **No access** from the dropdown menu.
 - iii. **Super user access**: Select **Yes** or **No**.
 - iv. **NFS version**: Select **All**, **NFSv3**, or **NFSv4**.
6. Select **Apply**.

Manage the SMB/CIFS share for a volume in NetApp Workload Factory

Managing a volume's SMB/CIFS share in Workload Factory includes SMB/CIFS share creation, determining the users and groups to give access to and the level of permissions to give them, and SMB/CIFS share deletion.

Before you begin

Before you begin, do the following:

- To manage SMB/CIFS shares, you must associate a link. [Learn how to associate an existing link or to create and associate a new link](#). After the link associates, return to this operation.
- Determine the users or groups to give access to and the level of permissions to give them.

Create an SMB/CIFS share for a volume

Follow the steps to create an SMB/CIFS share for a volume.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Volumes explorer**.
3. Select the actions menu of the volume to change the SMB share for.
4. Select **Advanced actions**, then **Manage SMB/CIFS shares**.
5. On the Manage SMB/CIFS shares page, select **Create SMB/CIFS share**.
6. In the Create SMB/CIFS share dialog, provide the following:
 - a. **Name**: Enter the name of the SMB/CIFS share.
 - b. **Path**: Either define the path using the default volume name or provide a share to an internal directory.

Valid path inputs for volume name, for example "avocado", are as follows:

- /avocado
- /avocado/folder
- /avocado/folder/subfolder
- /avocado/file-name

Valid path inputs for share name, for example "Server", are as follows:

- \\Server
- \\Server\Projects
- \\Server\Projects\Shared resources

- c. **Permissions**: Select Full control, Read/Write, Read, or No access, and then enter the users or groups separated by a semicolon (;). Users or groups are case sensitive and the user's domain must be included using the format "domain\username".

7. Select **Create**.

Change an SMB/CIFS share for a volume

Follow the steps to change the SMB/CIFS share settings for a volume.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. In the **FSx for ONTAP** tab, select the actions menu of the file system with the volume to update and then select **Manage**.
4. From the file system overview, select the **Volumes** tab.
5. From the Volumes tab, select the actions menu of the volume to change the SMB share for.
6. Select **Advanced actions**, then **Manage SMB/CIFS shares**.
7. On the Manage SMB/CIFS shares page, select **View and edit**.
8. Change the SMB/CIFS access permissions, or the users or groups to give permissions to.

Changes might cause current users or groups to lose access to the SMB/CIFS share.

9. Select **Apply** to save the changes.

Delete an SMB/CIFS share for a volume

Follow the steps to delete an SMB/CIFS share for a volume.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. In the **FSx for ONTAP** tab, select the actions menu of the file system with the volume to update and then select **Manage**.
4. From the file system overview, select the **Volumes** tab.
5. From the Volumes tab, select the actions menu of the volume to change the SMB share for.
6. Select **Advanced actions**, then **Manage SMB/CIFS shares**.
7. On the Manage SMB/CIFS shares page, select the actions menu of the SMB/CIFS share and then select **Delete**.

Deleting the SMB/CIFS share makes it unavailable and inaccessible to any users who want to mount it.

8. In the Delete SMB/CIFS share dialog, select **Delete** to confirm deletion.

Manage S3 access points

Create S3 access points for a volume in NetApp Workload Factory

Create and attach S3 access points for a volume in NetApp Workload Factory.

About this task

Amazon FSx for NetApp ONTAP lets NFS and SMB file systems access S3 data and connect to AWS services like Amazon Bedrock, SageMaker, Athena, AWS Glue, and more. You can connect AWS services to all of your object storage data.

Attach S3 access points to NFS and SMB volumes in an FSx for ONTAP file system so AWS services access files as if they are in an S3 bucket. When attaching the access point, define a unique ID, select the file access type (UNIX or Windows), and add a username for access authorization.

After you attach the S3 access point, it appears in the AWS Management Console with a unique alias. Use this alias as the S3 bucket name for AWS services, such as Amazon Bedrock, to access files in the FSx for ONTAP volume.

You can attach multiple S3 access points to a single FSx for ONTAP volume, each with a unique access level, to connect to multiple AWS services.

Before you begin

Complete the following requirements before you begin.

- You must have an existing volume with an S3 access point. [Create a volume with an S3 access point](#)
- You must [grant credentials with the operations and remediation permission policy](#) in Workload Factory to complete this task.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Volumes explorer**.
3. Select the actions menu for the volume you want to use, then select **Advanced actions**, and then **Manage S3 access points**.
4. On the Manage S3 access points page, select **Create and attach S3 access point**.
5. On the **Create and attach S3 access point** page, provide the following information:
 - **S3 access point name**: Enter the name of the S3 access point.
 - **User**: Select an existing user with access to the volume or create a new user.
 - **User type**: Select **UNIX** or **Windows** as the user type.
 - **Network configuration**: Select **Internet** or **Virtual private cloud (VPC)**. The type of network you choose determines whether the access point is accessible from the internet or restricted to a specific VPC. To use the Journal table feature, you must select **Internet** as the network configuration.
 - **Inventory table**: Optional. Requires the `getObject` permission. Enable the inventory table on the volume to generate metadata for all objects accessible to the S3 access point and incurs AWS S3 request costs. Refer to [Amazon S3 pricing documentation](#) for more information. The table updates every 24 hours.
 - **S3 access point tags**: Optionally, you can add up to 50 tags or remove tags.
6. Select **Create**.

Related information

Optionally, you can use the Journal table feature with S3 access points in Workload Factory. The Journal table infrastructure captures and stores audit logs of user access events and object operations across Amazon FSx for ONTAP volume access points. To enable the Journal table feature, you must set up the necessary AWS infrastructure, including AWS CloudTrail, AWS CloudWatch, AWS S3 Buckets, AWS CloudWatch log groups, and AWS Identity and Access Management (IAM) roles and policies. [Set up the journal table infrastructure for](#)

[NetApp Workload Factory](#).

Other related topics:

[Edit or delete S3 access points for a volume in NetApp Workload Factory](#)

Manage the S3 access points for a volume in NetApp Workload Factory

Update S3 access point details or delete S3 access points for a volume in NetApp Workload Factory.

You can also enable or disable the **Inventory table** for the access point. The inventory table generates metadata for all objects accessible to the S3 access point and incurs AWS S3 request costs. Refer to [Amazon S3 pricing documentation](#) for more information.

Edit the S3 access points for a volume

Edit the S3 access points for a volume in NetApp Workload Factory.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > Volumes explorer**.
3. Select the actions menu for the volume you want to edit S3 access points for, then select **Advanced actions**, and then **Manage S3 access points**.
4. From the **Manage S3 access points** screen, select the actions menu and then select **Edit access point**.
5. Make updates and then select **Apply**.

To set up the infrastructure for the journal table feature, you must select **Internet** as the network configuration for the S3 access point. [Learn how to set up the Journal table infrastructure for NetApp Workload Factory](#).

View S3 access point details for a volume

View S3 access point details like the alias, ARN, and S3 URI. With inventory and journal tables enabled, you can view metadata and audit logs of user access events and object operations across S3 access points.

About this task

When the inventory table is enabled on the volume, you can view access point, inventory table, and table bucket details of existing S3 access points attached to the volume. The system also provides a link to the inventory table in the AWS Management Console.

When the journal table is enabled, Workload Factory generates logs on access points in a table for your reference. The table includes CloudTrail events, trails (a record of AWS activities), CloudTrail filters, CloudWatch log groups, and more.

You can copy access point details for use in other applications.

Steps

1. Log in to the [NetApp Console](#).
- 2.

Select the menu , **Workloads > Storage > Volumes explorer**.

3. Select the actions menu for the volume to view details for, then select **Advanced actions**, and then **View details**.

Details appear under four tabs: **Access point**, **Inventory table**, **Journal table**, and **Table bucket**.

4. If the status for inventory or journal tables is enabled, but you're missing the `get:Object` permission, follow the instructions in the dialog to add the required permission.
5. When the journal table status is pending, you must provide the CloudTrail ARN to connect the journal table to the created infrastructure. Follow the instructions to deploy the CloudFormation stack using the AWS CLI or AWS Management Console and then return to the Workload Factory console to enter the CloudTrail ARN.

Delete the S3 access points for a volume in NetApp Workload Factory

Delete existing S3 access points from a volume in NetApp Workload Factory.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > Volumes explorer**.
3. Select the actions menu for the volume you want to manage S3 access points for, then select **Advanced actions**, and then **Manage S3 access points**.
4. Select the actions menu for the S3 access point to delete and then select **Delete**.
5. In the **Delete an S3 access point** dialog, select **Delete** to delete the S3 access point from the volume.

Workload Factory doesn't delete the inventory or journal table, but you can remove them manually from the AWS Management Console.

Related information

- [Set up the journal table infrastructure](#)

Set up the Journal table infrastructure for NetApp Workload Factory

Set up the Journal table infrastructure to capture and store audit logs of user access events and object operations across Amazon FSx for ONTAP volume access points. Several steps are necessary to set up the infrastructure for AWS services like AWS CloudTrail, AWS CloudWatch, AWS S3 Buckets, AWS CloudWatch log group, AWS Identity and Access Management (IAM), and AWS S3 Tables so that the log events travel through the pipeline correctly and are read by Workload Factory.

About this task

The Journal table feature captures S3 data-plane events (PutObject, GetObject, DeleteObject, etc.) for monitored FSx for ONTAP S3 access points. It uses a chain of AWS services deployed into your AWS account. When you set up the infrastructure correctly, it connects to the FSx for ONTAP volume access point and establishes the pipeline that captures user access and object operation audit events in the Journal table.

The following table lists the AWS services that are part of the infrastructure, their respective resource name

patterns, and the purpose of the service in the pipeline.

AWS service	Resource name pattern	Purpose
AWS CloudFormation	netapp-metadata-*	Deploys all infrastructure as a stack
AWS S3 bucket	netapp-metadata-cloudtrail-events-logs-{uuid}	Stores raw CloudTrail log files
AWS CloudTrail	netapp-metadata-journal-data-events-trail-{uuid}	Captures S3 data events for specific access points
AWS CloudWatch log group	netapp-metadata-journal-data-events-{uuid}	Receives CloudTrail events as structured log entries
IAM roles	- netapp-metadata-cloudtrail-cw-role-{uuid} - netapp-metadata-s3table-integration-role-{uuid}	Active integration
ObservabilityAdmin	S3TableIntegration	Bridges CloudWatch Logs into an S3 Tables table
S3 Tables	aws-cloudwatch bucket → logs.aws_cloudtrail__data	Stores structured, queryable CloudTrail events in Iceberg format

The {uuid} is a random 8-character identifier generated when the template is created.

Before you begin

To enable the Journal table feature, complete these steps:

- Have an existing volume with an S3 access point. [Create a volume with an S3 access point](#)
- Set network configuration to **Internet** for the S3 access point. [Edit network configuration for the S3 access point](#).
- [Grant the operations and remediation permissions](#) to your Workload Factory credentials.
- Add the following IAM policy permissions to the AWS account you use to run the CloudFormation deployment to set up the Journal table.

▼ IAM policy permissions for Journal table setup

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "CFNStack",
      "Effect": "Allow",
      "Action": [
        "cloudformation:CreateStack",
        "cloudformation:DescribeStacks",
        "cloudformation:DescribeStackEvents"
      ],
      "Resource": "arn:aws:cloudformation:*:*:stack/netapp-
```

```

metadata-*/*"
    },
    {
        "Sid": "StarResources",
        "Effect": "Allow",
        "Action": [
            "cloudformation:GetTemplateSummary",
            "cloudtrail:DescribeTrails",
            "logs:DescribeLogGroups",
            "logs:ListSourcesForS3TableIntegration",
            "observabilityadmin:CreateS3TableIntegration",
            "observabilityadmin:GetS3TableIntegration",
            "observabilityadmin:TagResource",
            "observabilityadmin:ListTagsForResource"
        ],
        "Resource": "*"
    },
    {
        "Sid": "S3Bucket",
        "Effect": "Allow",
        "Action": [
            "s3:CreateBucket",
            "s3:PutBucketPolicy",
            "s3:PutBucketTagging"
        ],
        "Resource": "arn:aws:s3:::netapp-metadata-*"
    },
    {
        "Sid": "IAMRoles",
        "Effect": "Allow",
        "Action": [
            "iam:CreateRole",
            "iam:PutRolePolicy",
            "iam:TagRole",
            "iam:GetRole"
        ],
        "Resource": "arn:aws:iam::*:role/netapp-metadata-*"
    },
    {
        "Sid": "PassRole",
        "Effect": "Allow",
        "Action": "iam:PassRole",
        "Resource": "arn:aws:iam::*:role/netapp-metadata-*",
        "Condition": {
            "StringEquals": {
                "iam:PassedToService": [

```

```

        "cloudtrail.amazonaws.com",
        "logs.amazonaws.com"
    ]
}
}
},
{
    "Sid": "CloudTrail",
    "Effect": "Allow",
    "Action": [
        "cloudtrail:CreateTrail",
        "cloudtrail:StartLogging",
        "cloudtrail:AddTags",
        "cloudtrail:PutEventSelectors"
    ],
    "Resource": "arn:aws:cloudtrail:*:*:trail/netapp-metadata-
* "
},
{
    "Sid": "CWLogGroup",
    "Effect": "Allow",
    "Action": [
        "logs:CreateLogGroup",
        "logs>DeleteLogGroup",
        "logs:PutRetentionPolicy",
        "logs:TagResource",
        "logs:AssociateSourceToS3TableIntegration"
    ],
    "Resource": "arn:aws:logs:*:*:log-group/netapp-metadata-*"
},
{
    "Sid": "S3Table",
    "Effect": "Allow",
    "Action": [
        "s3tables:CreateTableBucket",
        "s3tables:PutTableBucketEncryption",
        "s3tables:PutTableBucketPolicy"
    ],
    "Resource": "arn:aws:s3tables:*:*:bucket/aws-cloudwatch"
}
]
}

```

Set up the journal table infrastructure

Set up the infrastructure to capture AWS service events from the S3 access point in the journal table.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From **FSx for ONTAP**, select the actions menu of the file system with the volume to update, then select **Manage**.
4. From the file system overview, select the **Volumes** tab.
5. From the Volumes tab, select the actions menu for the volume you want to manage S3 access points for, then select **Advanced actions**, and then **Manage S3 access points**.
6. From the **Manage S3 access points** screen, select the actions menu and then select **Edit access point**.
7. In the **Edit S3 access point** dialog, ensure that the network configuration is set to **Internet**.
8. Follow the instructions in the dialog to set up the infrastructure for the journal table feature.
9. Download the CloudFormation template.
10. Deploy the CloudFormation stack in your AWS account.
 - a. Save the template JSON to a file.
 - b. Deploy the template using the AWS CLI or AWS Management Console.
 - c. Wait for the stack to reach `CREATE_COMPLETE` status.
 - d. Get the CloudTrail ARN from the stack outputs.
11. Return to the Workload Factory console and back to the volume to manage S3 access points for.
12. Select **View details** from the volume actions menu.
13. In the **Journal table** tab, enter the CloudTrail ARN.
14. Select **Apply**.

If any of the steps fail, [troubleshoot the infrastructure setup for the journal table](#) or contact NetApp support for assistance.

Troubleshoot the infrastructure setup for the journal table

You can use the AWS Management Console or the AWS CLI to troubleshoot failures during the CloudFormation stack deployment and the resources it creates.

After resolving the issue, submit the CloudTrail ARN again to restart journal table setup.

Common CloudFormation deployment failures

Insufficient IAM permissions

The deploying role/user needs a specific set of permissions. Refer to [Before you begin](#) for the full policy.

If the stack fails with an `AccessDenied` or `InsufficientPermission` error, check the stack events and look for the missing permission in the `ResourceStatusReason` field.

CloudTrail limit

AWS enforces a default limit of 5 trails per region. If the account already has 5 trails, the `FsxDataEventTrail` resource will fail with: `Maximum number of trails (5) exceeded`. You can check the

number of Trails on a regional level with the following command:

```
aws cloudtrail describe-trails \  
  --no-include-shadow-trails \  
  --region <region> \  
  --query "length(trailList) "
```

Resolution options

- Option 1: Delete an unused trail in the region to make room.
- Option 2: Use an existing trail. Remove the `FsxDataEventTrail`, `CloudTrailBucket`, `CloudTrailBucketPolicy`, and `CloudTrailToCloudWatchRole` resources from the template before deploying the CloudFormation stack again. Then pass the ARN of your existing trail during the initiation step. The existing trail must have a CloudWatch Log Group configured, an `S3TableIntegration` associated with the log group, and be logging data events.

S3 Table integration already exists

If the account already has an `S3TableIntegration` for the `aws_cloudtrail` data source, the `LogsToS3TableIntegration` resource will fail.

Resolution

Remove the `LogsToS3TableIntegration` and `S3TableIntegrationRole` resources from the template before deploying the CloudFormation stack again. The system automatically uses the existing integration as long as you configure it for `aws_cloudtrail` data events.

To check for an existing integration:

```
aws observabilityadmin list-s3-table-integrations --region <your-region>
```

S3 bucket name already exists

The bucket name `netapp-metadata-cloudtrail-events-logs-{uuid}` is globally unique. If it collides, re-request the template to get a new UUID.

IAM role already exists

If a previous partial deployment left behind IAM roles with the `netapp-metadata-*` name pattern, the stack will fail on role creation. Delete the orphaned roles first:

```
aws iam delete-role-policy \  
  --role-name netapp-metadata-cloudtrail-cw-role-<uuid> \  
  --policy-name <policy-name>  
aws iam delete-role \  
  --role-name netapp-metadata-cloudtrail-cw-role-<uuid>
```

Failures after Journal table enablement

After submitting the CloudTrail ARN, Workload Factory validates the entire resource pipeline by sending a *seed*, or test, event automatically. If successful, the seed event arrives in the S3 Tables table. The test takes approximately 10 minutes.

If the test validation fails, you might get one of the following error messages:

Error message	Meaning
Table aws_cloudtrail__data was not created in {bucket}. Verify s3table permissions.	S3TableIntegration did not create the CloudWatch-managed table. The pipeline between CloudWatch Logs and S3 Tables is broken.
Table exists, but the journal seed event does not appear. Verify CloudTrail and CloudWatch permissions.	The table exists but the specific seed event never arrived. Pipeline is broken between CloudTrail and the S3 Tables table.
Failed to initiate journal setup. ...	An error occurred during the background seed/poll flow. Check the trailing message for details.

Resolution steps

When the journal reaches FAILED, trace the seed event forward through the pipeline stages to identify exactly where it stopped. Each step maps to a specific AWS resource created by the template.

1. Check the CloudTrail S3 Bucket.

The trail writes raw event logs to the S3 bucket `netapp-metadata-cloudtrail-events-logs`. Look for recent log files.

If no log files exist, then CloudTrail is not capturing events. Check the following:

- The trail is logging (`IsLogging: true`)
- The advanced event selectors include the correct access point ARN
- The advanced event selectors include the filters for `eventCategory = Data` and `resources.type = AWS::S3::AccessPoint`

2. Check the CloudWatch Log Group.

The trail also delivers events to the CloudWatch Log Group. The log group name starts with `netapp-metadata-journal-data-events-<uuid>`.

- If the log group is empty, then CloudTrail is not delivering events to CloudWatch. Check that the `CloudTrailToCloudWatchRole` IAM role exists and has `logs:CreateLogStream` and `logs:PutLogEvents` permissions, and that the trail is configured with the correct `CloudWatchLogsLogGroupArn` and `CloudWatchLogsRoleArn`.
- If the seed event appears in the log group, then the problem is downstream — proceed to Step 3.

3. Check the S3 Tables table (`aws-cloudwatch`).

The `S3TableIntegration` automatically creates a table bucket called `aws-cloudwatch` and populates a table at `logs.aws_cloudtrail__data`. This table is only created after the first event flows through.

- If the `aws-cloudwatch` table bucket does not exist, then the `S3TableIntegrationRole` is missing

permissions. It needs `s3tables:CreateTableBucket`, `s3tables:PutTableBucketEncryption`, and `s3tables:PutTableBucketPolicy` — all scoped to `arn:aws:s3tables:*:*:bucket/aws-cloudwatch`.

- If the table bucket exists but `logs.aws_cloudtrail__data` does not, then the integration is not routing events. The integration must show `Status: ACTIVE` and include `aws_cloudtrail` as a log source.
- If the table exists but the seed event is not in it, then the event may still be in transit. S3 Tables ingestion has some latency. Wait a few more minutes. If it still does not appear after 15-20 minutes, the integration may be broken.

4. Query the seed event directly.

- Open the S3 Tables in the AWS Management Console.
- Navigate to the `aws-cloudwatch` table bucket → `aws_cloudtrail__data` table, and use the **Preview** button to run a quick query directly in the browser.
- If the event is present in the table but the journal still shows `FAILED`, then the polling window may have expired before the event arrived.

After resolving the issue, return to the Workload Factory console. Retry [initiating the journal table setup](#) by submitting the Trail ARN again.

- If setup continues to fail, contact NetApp support for assistance.

Permissions reference for the journal table setup

The IAM role that deploys the CloudFormation stack to enable the Journal table feature needs the following permissions. Refer to [Before you begin](#) for a copiable JSON policy with the required permissions.

Stack Operations

Permission	Resource	Why
<code>cloudformation:CreateStack</code>	<code>arn:aws:cloudformation:*:*:stack/netapp-metadata-/*</code>	Create the stack
<code>cloudformation:DescribeStacks</code>	<code>arn:aws:cloudformation:*:*:stack/netapp-metadata-/*</code>	Monitor stack status
<code>cloudformation:DescribeStackEvents</code>	<code>arn:aws:cloudformation:*:*:stack/netapp-metadata-/*</code>	Diagnose resource-level failures <code>cloudformation:GetTemplateSummary</code> * Pre-flight template validation

CloudTrail

Permission	Resource	Why
<code>cloudtrail:CreateTrail</code>	<code>arn:aws:cloudtrail:*:*:trail/netapp-metadata-/*</code>	Create the trail
<code>cloudtrail:StartLogging</code>	<code>arn:aws:cloudtrail:*:*:trail/netapp-metadata-/*</code>	Enable logging
<code>cloudtrail:AddTags</code>	<code>arn:aws:cloudtrail:*:*:trail/netapp-metadata-/*</code>	Apply identification tag

Permission	Resource	Why
cloudtrail:PutEventSelectors	arn:aws:cloudtrail:*:*:trail/netapp-metadata-*	Configure data event capture
cloudtrail:DescribeTrails	*	Resolve trail ARN for stack output

S3

Permission	Resource	Why
s3:CreateBucket	arn:aws:s3:::netapp-metadata-*	Create the CloudTrail log bucket
s3:PutBucketPolicy	arn:aws:s3:::netapp-metadata-*	Allow CloudTrail to write logs
s3:PutBucketTagging	arn:aws:s3:::netapp-metadata-*	Apply identification tag

IAM

Permission	Resource	Why
iam:CreateRole	arn:aws:iam:*:*:role/netapp-metadata-*	Create both IAM roles
iam:PutRolePolicy	arn:aws:iam:*:*:role/netapp-metadata-*	Attach inline policies
iam:TagRole	arn:aws:iam:*:*:role/netapp-metadata-*	Apply identification tag
iam:GetRole	arn:aws:iam:*:*:role/netapp-metadata-*	Confirm role is active
iam:PassRole	arn:aws:iam:*:*:role/netapp-metadata-* (condition: PassedToService = cloudtrail.amazonaws.com, logs.amazonaws.com)	Pass roles to CloudTrail and CloudWatch Logs

CloudWatch Logs

Permission	Resource	Why
logs:CreateLogGroup	arn:aws:logs:*:*:log-group:netapp-metadata-*	Create the log group
logs>DeleteLogGroup	arn:aws:logs:*:*:log-group:netapp-metadata-*	Clean up log group if creation failed
logs:PutRetentionPolicy	arn:aws:logs:*:*:log-group:netapp-metadata-*	Set 30-day retention
logs:TagResource	arn:aws:logs:*:*:log-group:netapp-metadata-*	Apply identification tag
logs:AssociateSourceToS3TableIntegration	arn:aws:logs:*:*:log-group:netapp-metadata-*	Link CloudTrail source to S3 Tables
logs:DescribeLogGroups	*	Check log group existence

Permission	Resource	Why
logs:ListSourcesForS3TableIntegration	*	Confirm integration association

ObservabilityAdmin

Permission	Resource	Why
observabilityadmin:CreateS3TableIntegration	*	Create the CloudWatch → S3 Tables bridge
observabilityadmin:GetS3TableIntegration	*	Confirm integration is active
observabilityadmin:TagResource	*	Apply identification tag
observabilityadmin:ListTagsForResource	*	Drift detection

S3 Tables

Permission	Resource	Why
s3tables:CreateTableBucket	arn:aws:s3tables:*:*:bucket/aws-cloudwatch	Create the S3 Tables bucket (via integration role)
s3tables:PutTableBucketEncryption	arn:aws:s3tables:*:*:bucket/aws-cloudwatch	Set AES256 encryption
s3tables:PutTableBucketPolicy	arn:aws:s3tables:*:*:bucket/aws-cloudwatch	Allow CloudWatch Logs access

Split a cloned volume in NetApp Workload Factory

Split a cloned FlexVol volume from its parent volume to make the clone a normal read/write FlexVol volume in NetApp Workload Factory.

Data is accessible on the clone and the parent during the split. The split process only updates metadata and requires minimal IO. No data blocks are copied.

About this task

The clone splitting operation involves the following:

- New snapshot copies of the FlexClone volume cannot be created during the split operation.
- A FlexClone volume cannot be split from the parent volume if it belongs to a data protection relationship.
- If you take the FlexClone volume offline while splitting is in progress, the split operation is suspended; when you bring the FlexClone volume back online, the splitting operation resumes.
- After the split, both the parent FlexVol volume and the clone require the full space allocation determined by their volume guarantees.
- After a FlexClone volume is split from its parent the two cannot be rejoined.

Before you begin

Consider the following before you split a cloned volume:

- To split a cloned volume, you must associate a link. [Learn how to associate an existing link or to create and associate a new link](#). After the link associates, return to this operation.
- The FlexClone volume must be online when the split operation begins.
- The parent volume must be online for the split to succeed.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Volumes explorer**.
3. Select the actions menu of the volume with the cloned volume to split.
4. Select **Data protection actions**, then **Split cloned volume**.
5. In the Split volume dialog, select **Delete**.

Result

The volume clone is split and appears in the Volumes tab.

Delete a volume in NetApp Workload Factory

Delete a volume in your FSx for ONTAP file system that is no longer required and to free up space. This operation is irreversible.

Before you begin

Consider the following before deleting a volume:

- Replication relationships: You must [delete all existing replication relationships](#) for this volume before deleting the volume so that no broken relationships remain.
- Local snapshots: All snapshots associated with this FSx for ONTAP file system will be permanently deleted.
- Volume backups: Volume backup copies will remain and you can still use them.
- Immutable files and snapshots: Volumes containing immutable files and snapshots cannot be deleted until the retention period ends.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Volumes explorer**.
3. Select the actions menu of the volume to delete.
4. Select **Basic actions** then **Delete volume**.
5. In the Delete volume dialog, do the following:
 - a. Optionally, select **Back up the volume** to back up the volume before deletion.

The backup will remain in the file system until you manually delete it.

- b. Select **Continue**.

- c. Type “delete” to delete the volume.
- d. Select **Delete** to confirm.

Copyright information

Copyright © 2026 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.