



Protect your data

Amazon FSx for NetApp ONTAP

NetApp
February 02, 2026

This PDF was generated from <https://docs.netapp.com/us-en/workload-fsx-ontap/data-protection-overview.html> on February 02, 2026. Always check docs.netapp.com for the latest.

Table of Contents

| | |
|---|----|
| Protect your data | 1 |
| Types of data protection in NetApp Workload Factory | 1 |
| Types of data protection | 1 |
| Best practices for protecting your workload data | 2 |
| Protect your workload data with snapshots | 2 |
| Protect your workload data with NetApp Autonomous Ransomware Protection with AI | 2 |
| Protect your workload data with volume replication | 2 |
| Protect your workload data with backups | 3 |
| Recommendations for protecting your workload data | 3 |
| Use snapshots | 3 |
| Create a manual snapshot of an FSx for ONTAP volume | 3 |
| Create a snapshot policy for storage VMs in Workload Factory | 4 |
| Restore a volume from a snapshot in Workload Factory | 5 |
| Use backups to object storage | 6 |
| Create a manual backup of a volume in NetApp Workload Factory | 6 |
| Restore a volume from a backup in NetApp Workload Factory | 7 |
| Use replication | 8 |
| Create a replication relationship in NetApp Workload Factory | 8 |
| Initialize a replication relationship in NetApp Workload Factory | 11 |
| Protect your data with NetApp Autonomous Ransomware Protection with AI | 11 |
| Enable ARP/AI for a file system or a volume | 12 |
| Validate ransomware attacks | 13 |
| Recover data after a ransomware attack | 14 |
| Clone a volume in NetApp Workload Factory | 14 |
| Use on-premises ONTAP cluster data in NetApp Workload Factory | 15 |
| Discover an on-premises ONTAP cluster | 16 |
| Replicate volume data from an on-premises ONTAP cluster | 17 |
| Remove an on-premises ONTAP cluster from NetApp Workload Factory | 18 |
| Protect your data with a cyber vault | 19 |

Protect your data

Types of data protection in NetApp Workload Factory

FSx for ONTAP supports snapshots, NetApp Autonomous Ransomware Protection with AI, replication, and backups for data protection. We recommend that you use a combination of data protection types to prepare for the inevitable and safeguard your data.

Types of data protection

Data protection for your workloads helps ensure that you can recover from any data loss at any time. Learn about the types of data protection before you select the features you'll use.

Snapshots

A snapshot creates a read-only, point-in-time image of a volume within the source volume as a snapshot copy. You can use the snapshot copy to recover individual files, or to restore the entire contents of a volume.

Snapshots are the basis of all the backup methods. The snapshot copy that is created on your volume is used to keep the replicated volume and backup file synchronized with changes made to the source volume.

NetApp Autonomous Ransomware Protection with AI

NetApp Autonomous Ransomware Protection with AI (ARP/AI) uses workload analysis in NAS (NFS/SMB) environments to detect and warn about abnormal activity that might be a ransomware attack. When an attack is suspected, ARP/AI also creates new, immutable snapshots in addition to the existing protection provided by scheduled snapshots.

Replication

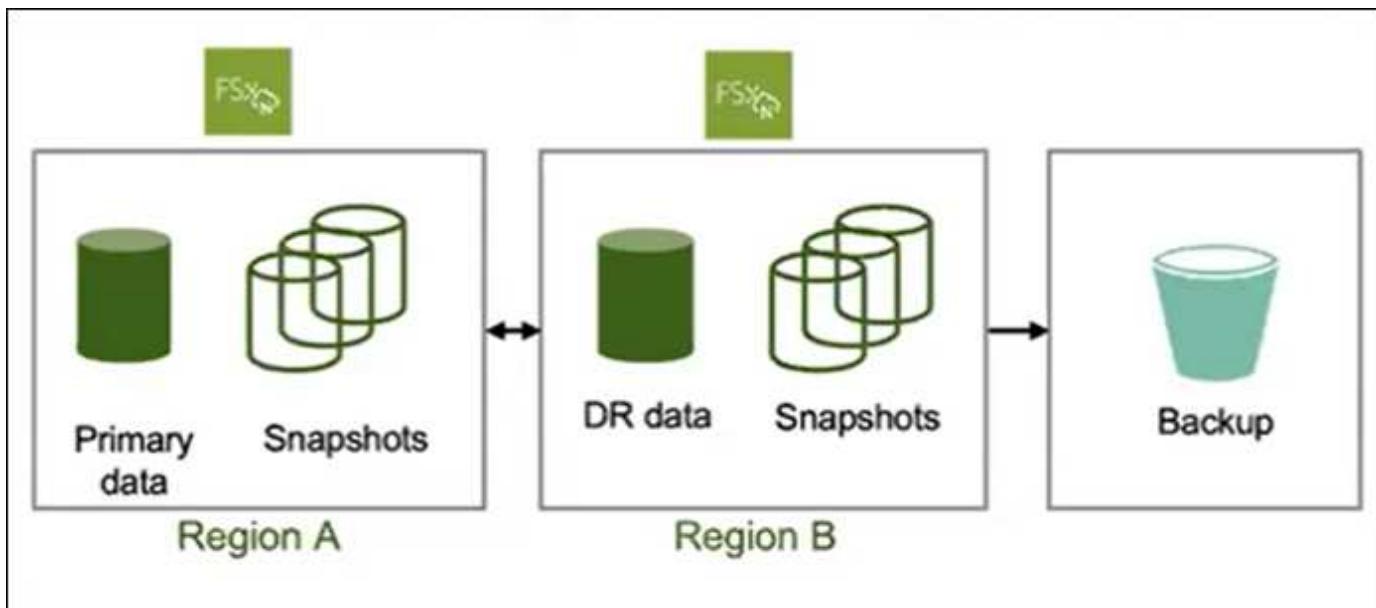
Replication creates a secondary copy of your data on another FSx for ONTAP file system and continually updates the secondary data. Your data is kept current and remains available whenever you need it, such as for disaster recovery.

You can choose to create both replicated volumes on another FSx for ONTAP file system and backup files in the cloud. Or you can choose just to create replicated volumes or backup files - it's your choice.

Backups

You can create backups of your data to the cloud for protection and for long-term retention purposes. If necessary, you can restore a volume, folder, or individual files from the backup to the same, or different, working file system.

The following diagram shows a visual representation of data protection for FSx for ONTAP storage using snapshots, replication across regions, and backup to object storage.



Best practices for protecting your workload data

FSx for ONTAP offers multiple data protection options which can be combined to achieve your selected recovery point and time objectives. For the best possible protection, we recommend that you use both volume snapshots and volume backups.

A recovery point objective (RPO) describes how recent the latest copy of your data is guaranteed to be, which depends on how frequently the copies are made. A recovery time objective (RTO) defines how long it takes to restore your data.

Protect your workload data with snapshots

Snapshots are virtual point-in-time versions of a volume that are taken on a scheduled basis. You can access snapshots using standard file system commands. Snapshots provide an RPO of as little as one hour. RTO depends on the amount of data to restore and is primarily limited by the volume throughput limit. Snapshots also allow users to restore specific files and directories, which decreases RTO even further. Snapshots only consume additional volume space for changes made to the volume.

Protect your workload data with NetApp Autonomous Ransomware Protection with AI

NetApp Autonomous Ransomware Protection with AI (ARP/AI) acts as an important additional layer of defense if anti-virus software has failed to detect an intrusion. Setting an ARP/AI policy enables it for all storage VMs and all existing and newly created volumes. Once enabled, ARP/AI detects and protects all volumes and storage VMs. If a file extension is flagged as abnormal, you should evaluate the alert.

Protect your workload data with volume replication

Volume replication creates a copy of the latest data of a volume including all its snapshots in a different region. If you cannot afford multi-hour RTOs of a full volume restore operation from a volume backup, consider performing a volume replication. While volume replication makes sure recent data is available in a different region for you to use, you need to adjust your clients to use the volume in the other region.

Protect your workload data with backups

Volume backups provide independent point-in-time copies of your volume. They can be used to store old backups and provide the necessary second copy of your data. Daily, weekly, and monthly backup schedules allow for RPOs starting at one day. Volume backups can only be restored as a whole. Creating a volume from a backup (RTO) can take hours to many days, depending on the size of the backup.

Recommendations for protecting your workload data

Consider the following recommendations for protecting your workload data.

- Use volume replication for disaster recovery: if your application requires a low RTO, consider using volume replication to replicate your data to another region.
- Use volume backups in conjunction with snapshots: using the two features together ensures that you're able to restore your files from snapshots and perform full restores in case of volume loss using backups.
- Define a volume backup policy: make sure that the backup policy satisfies your company requirements for backup age and frequency. We recommend keeping a minimum of two daily backups for each volume.
- Define a snapshot schedule: older snapshots are less likely to be used to restore data. We recommend that you define a snapshot schedule that takes into consideration the diminishing returns of keeping older snapshots against the cost for additional snapshot capacity.
- Enable an ARP/AI policy for your file system or individual volumes to add an additional layer of protection to protect your data from ransomware attacks.

Use snapshots

Create a manual snapshot of an FSx for ONTAP volume

Create a manual snapshot of an FSx for ONTAP volume in NetApp Workload Factory. Snapshots are point-in-time versions of your volume's content.

Snapshots are resources of volumes and are instant captures of your data that consume space only for modified data. Because data changes over time, snapshots usually consume more space as they get older.

FSx for ONTAP volumes use just-in-time copy-on-write so that any unmodified files in snapshots don't consume any of the volume's capacity.



Snapshots aren't copies of your data. If you want to make copies of your data, consider using the FSx for ONTAP backups or volume replication features.

Before you begin

You must associate a link to create a manual snapshot of a volume. [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 using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the actions menu of the file system that contains the volume to create a snapshot for and then select **Manage**.

5. In the file system overview, select the **Volumes** tab.
6. From the **Volumes** tab, select the actions menu for the volume to protect with snapshots.
7. Select **Data protection actions** and then **Manage snapshots**.
8. From the Manage snapshots page, select **Create snapshot**.
9. In the Create a snapshot dialog, do the following:
 - a. Enter a snapshot name in the **Snapshot name** field.
 - b. Optionally, select a label or create a new label.
 - c. Set the **Retention period** as a number of hours, days, months, or years.
 - d. Optional: **Make this snapshot immutable** to prevent the snapshot from being deleted during the retention period.

Accept the statement regarding immutable snapshots.

10. Select **Create**.

Create a snapshot policy for storage VMs in Workload Factory

Create a custom snapshot policy for storage VMs in Workload Factory to manage snapshot creation and retention. A snapshot policy defines how the system creates snapshots for a storage VM. You can create a snapshot policy for a storage VM in an FSx for ONTAP file system. You can also share the policy across multiple storage VMs.

About this task

You can create a custom snapshot policy that differs from the three built-in snapshot policies for FSx for ONTAP:

- default
- default-1weekly
- none

By default, every volume is associated with the file system's `default` snapshot policy. We recommend using this policy for most workloads.

Customizing a policy lets you specify when to create snapshots, how many copies to retain, and how to name them.

Before you begin

- Once a snapshot policy is created, its association with the storage VM(s) cannot be modified, but you can always add or remove the policy from volumes.
- Consider the following about snapshot capacity before you use snapshots:
 - For most datasets, an additional capacity of 20% is enough to keep snapshots for up to four weeks. As data gets older, its use for restorations becomes less likely.
 - Overwriting all the data in a snapshot consumes significant volume capacity, which factors into provisioning volume capacity.
- To create a custom snapshot policy, 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 using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the actions menu of the file system with the volume and then select **Manage**.
5. In the file system overview, select the **Storage VMs** tab.
6. From the **Storage VMs** tab, select the actions menu for the volume to protect with scheduled snapshots, then **Advanced actions**, and then **Manage snapshot policies**.
7. On the Snapshot policy management page, select **Create snapshot policy**.
8. In the **Snapshot policy name** field, enter a name for the snapshot policy.
9. Optionally, enter a description for the snapshot policy.
10. Under **Policy schedule and copies**, select when to create snapshots. For example, every minute or hourly.

You can select more than one frequency.

11. Under **Number of copies**, enter the number of copies to retain.

The maximum number of copies is 1,023.

12. Optional: Under **Naming convention**, enter a **Prefix** for the policy.

13. **Retention label** is automatically populated.

This label refers to the SnapMirror, or replication label that is used to select only specified snapshots for replication from the source to the target file system.

14. Optional: Enable **Immutable snapshots** for any schedules you need, set the **Retention period** for each schedule, and accept the statement to continue.

Enabling immutable snapshots locks all snapshots in this snapshot policy to prevent the snapshots from being deleted during the retention period.

15. **Share across storage VMs**: Enabled by default. When enabled, the snapshot policy is shared across all storage VMs in the file system. Disable to create a snapshot policy for a single storage VM.

16. Select **Create**.

Restore a volume from a snapshot in Workload Factory

In Workload Factory, you can restore data from a snapshot to an existing volume or to a new volume. The restore operation enables point-in-time recovery when a volume contains deleted or corrupted files.

About this task

You have the option to restore data from a snapshot to an existing volume or to a new volume.

The creation of a new volume from a snapshot makes a copy of an entire volume within a few seconds independent of volume size. The newly created copy represents a new volume.

Before you begin

Consider the following limitations before you create a volume from a snapshot:

- You can only restore a volume from a snapshot if you have an existing snapshot copy of the volume.
- Changes to permission models: If you use this operation to switch the network-attached storage (NAS) protocol type, it might also switch the permission model that the security style provides. You might experience file access permission issues, which you can only fix manually with administrator access using the NAS client tools for permissions setting.

Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the actions menu of the file system with the volume and then select **Manage**.
5. In the file system overview, select the **Volumes** tab.
6. From the **Volumes** tab, select the actions menu for the volume to restore from a snapshot.
7. Select **Data protection actions** and then **Manage snapshots**.
8. From the Manage snapshots page, select the actions menu for the snapshot to restore, and then select **Restore**.
9. In the Restore volume from a snapshot dialog, select from the following options:
 - Toggle to select **Restore as a new volume**.
In the **Restored volume name** field, enter a unique name for the volume to restore.
 - Restore data from a snapshot to an existing volume. This operation permanently deletes any data that was modified after the snapshot creation time.Accept the statement to proceed.
10. Select **Restore**.

Use backups to object storage

Create a manual backup of a volume in NetApp Workload Factory

Create a manual backup of a volume outside regularly scheduled backups in NetApp Workload Factory.

About this task

FSx for ONTAP backups are per volume, so each backup contains only the data in a particular volume.

FSx for ONTAP backups are incremental which means that only the data on the volume that has changed after your most recent backup is saved. This minimizes the time required to create the backup and the storage required for the backup, which saves on storage costs by not duplicating data.

Before you begin

To take backups of your volumes, both your volume and your file system must have enough available SSD storage capacity to store the backup snapshot. When taking a backup snapshot, the additional storage capacity consumed by the snapshot cannot cause the volume to exceed 98% SSD storage utilization. If this

happens, the backup will fail.

Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the actions menu of the file system with the volume and then select **Manage**.
5. In the file system overview, select the **Volumes** tab.
6. From the **Volumes** tab, select the actions menu for the volume to back up.
7. Select **Data protection actions**, **FSx for ONTAP backup**, and then **Manual backup**.
8. In the Manual backup dialog, enter a name for the backup.
9. Select **Back up**.

Restore a volume from a backup in NetApp Workload Factory

In NetApp Workload Factory, you can restore a volume from a backup to any FSx for ONTAP file system in your AWS account.

Workload factory determines if you have enough capacity for the restore and can automatically add SSD storage tier capacity if you don't.

Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the actions menu of the file system with the volume and then select **Manage**.
5. In the file system overview, select the **Volumes** tab.
6. From the **Volumes** tab, select the actions menu for the volume to restore from a backup.
7. Select **Data protection actions**, **FSx for ONTAP backup**, and then **Restore from a backup**.
8. In the Restore from a backup dialog, provide the following:
 - a. **Target file system**: Select the target file system from the dropdown menu.
 - b. **Target storage VM**: Select the target storage VM from the dropdown menu.
 - c. **Backup name**: Select the backup name from the dropdown menu.
 - d. **Restored volume name**: Enter the restored volume name.
9. Verify file system capacity for the restore operation.

When file system capacity is limited, the following might occur:

- The restore can push used capacity over the threshold you specified. You can complete the restore operation. Consider [manually adding SSD storage tier capacity](#) or selecting for Workload Factory to automatically add SSD storage tier capacity.
- The restore requires additional SSD capacity. You must select for Workload Factory to automatically add SSD storage tier capacity to proceed.

10. Select **Restore**.

Use replication

Create a replication relationship in NetApp Workload Factory

Create a replication relationship for an FSx for ONTAP file system in NetApp Workload Factory to avoid data loss in case of an unforeseen disaster. Replication is supported between two FSx for ONTAP file systems, and between Cloud Volumes ONTAP or an on-premises ONTAP system and an FSx for ONTAP file system.

About this task

Replication adds protection against data loss if the region where your data resides experiences a disaster.

This operation creates a replication relationship for source volumes in an FSx for ONTAP file system, on-premises ONTAP system, or Cloud Volumes ONTAP system.

Replicated volumes in the target file system are data protection (DP) volumes and follow the naming format: {OriginalVolumeName}_copy.

When you replicate a source volume with immutable files, the target volume and file system stay locked until the retention period of the immutable files in the source volume ends. The immutable files feature is available when you [create a volume](#) for an FSx for ONTAP file system.

- Replication isn't supported for block volumes using iSCSI or NVMe protocols.
- You can replicate one source (read/write) volume or one data protection (DP) volume. Cascading replication is supported, but a third hop isn't. Learn more about [cascading replication](#).

Before you begin

Consider the following before you begin.

- You must have one FSx for ONTAP file system to use for the target in the replication relationship.
- The FSx for ONTAP file system you use for the replication relationship must have an associated link. [Learn how to associate an existing link or to create and associate a new link](#). After you associate the link, return to this operation.
- For replication from an on-premises ONTAP system to an FSx for ONTAP file system, make sure you have discovered the on-premises ONTAP system.

Follow these steps to replicate specific or all volumes in a file system.

Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the actions menu of the file system that contains the volume(s) to replicate and then select **Manage**.
5. Either replicate all volumes in a file system or replicate select volumes.

- To replicate all volumes in a file system: From the file system overview, select **Replicate data**.
- To replicate select volumes: From the file system overview, select the **Volumes** tab.

In the Volumes table, select one or more volumes and then select **Replicate data**.

6. On the Replicate data page, under Replication target, provide the following:

- Use case:** Select one of the following use cases for the replication. Depending on the selected use case, Workload Factory fills in the form with recommended values in accordance with best practices. You can accept the recommended values or make changes as you complete the form.
 - Migration: transfers your data to the target FSx for ONTAP file system
 - Hot disaster recovery: ensures high availability and rapid disaster recovery for critical workloads
 - Cold or archive disaster recovery:
 - Cold disaster recovery: uses longer recovery time objectives (RTO) and recovery point objects (RPO) to lower costs
 - Archive: replicates data for long-term storage and compliance
 - Other

Additionally, the use case selection determines the replication policy, or SnapMirror policy (ONTAP). The terms used to describe replication policies come from [ONTAP 9 documentation](#).

- For migration and other, the replication policy is called *MirrorAll/Snapshots*. *MirrorAll/Snapshots* is an asynchronous policy for mirroring all snapshots and the latest active file system.
- For hot, cold, or archive disaster recovery, the replication policy is called *MirrorAndVault*. *MirrorAndVault* is an asynchronous and vault policy for mirroring the latest active file system and daily and weekly snapshots.

For all use cases, if you enable snapshots for long-term retention, the default replication policy is *MirrorAndVault*.

- FSx for ONTAP file system:** Select credentials, region, and FSx for ONTAP file system name for the target FSx for ONTAP file system.
- Storage VM name:** Select the storage VM from the dropdown menu. The storage VM you select is the target for all selected volumes in this replication relationship.
- Volume name:** The target volume name is generated automatically with the following format {OriginalVolumeName}_copy. You can use the auto-generated volume name or enter another volume name.
- Tiering policy:** Select the tiering policy for the data stored in the target volume. The tiering policy defaults to the recommended tiering policy for the use case you selected.

Balanced (Auto) is the default tiering policy when creating a volume using the Workload Factory console. For more information about volume tiering policies, refer to [Volume storage capacity](#) in AWS FSx for NetApp ONTAP documentation. Note that Workload factory uses use-case based names in the Workload Factory console for tiering policies and includes FSx for ONTAP tiering policy names in parentheses.

If you selected the migration use case, Workload Factory automatically selects to copy the tiering policy of source volume to the target volume. You can deselect to copy the tiering policy and select a tiering policy which applies to the volume selected for replication.

f. **Max transfer rate:** Select **Limited** and enter the max transfer limit in MB/s. Alternatively, select **Unlimited**.

Without a limit, network and application performance may decline. Alternatively, we recommend an unlimited transfer rate for FSx for ONTAP file systems for critical workloads, for example, those that are used primarily for disaster recovery.

7. Under Replication settings, provide the following:

- Replication interval:** Select the frequency that snapshots are transferred from the source volume to the target volume.
- Long-term retention:** Optionally, enable snapshots for long-term retention. Long-term retention enables business services to continue operating even through a complete site failure, supporting applications to fail over transparently using a secondary copy.

Replications without long-term retention use the *MirrorAllSnapshots* policy. Enabling long-term retention assigns the *MirrorAndVault* policy to the replication.

If you enable long-term retention, then select an existing policy or create a new policy to define the snapshots to replicate and the number to retain.



Matching source and target labels are required for long-term retention. If desired, Workload factory can create missing labels for you.

- **Choose an existing policy:** select an existing policy from the dropdown menu.
- **Create a new policy:** enter a **policy name**.

- Immutable snapshots:** Optional. Select **Enable immutable snapshots** to prevent snapshots taken in this policy from being deleted during the retention period.
 - Set the **Retention period** in number of hours, days, months, or years.
 - **Snapshot policies:** In the table, select the snapshot policy frequency and the number of copies to retain. You can select more than one snapshot policy.
- S3 access point:** Optionally, attach an S3 access point to access FSx for ONTAP file system data residing on NFS or SMB/CIFS volumes via AWS S3 APIs. Only the file access type is supported. Providing the following details:
 - **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.
 - **Enable metadata:** Enabling metadata creates an S3 table containing all objects accessible by the S3 access point, which you can use for auditing, governance, automatic, analysis, and optimization. Enabling metadata incurs additional AWS costs. Refer to [Amazon S3 pricing documentation](#) for more information.
- S3 access point tags:** Optionally, you can add up to 50 tags.

8. Select **Create**.

Result

The replication relationship appears in the **Replication relationships** tab in the target FSx for ONTAP file system.

Initialize a replication relationship in NetApp Workload Factory

Initialize a replication relationship between source and target volumes to transfer the snapshot and all data blocks in NetApp Workload Factory.

About this task

Initialization performs a *baseline* transfer: it makes a snapshot of the source volume, then transfers the snapshot and all the data blocks it references to the target volume.

Before you begin

Consider when you choose to complete this operation. Initialization can be time-consuming. You might want to run the baseline transfer in off-peak hours.

Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the actions menu of the file system to update and then select **Manage**.
5. From the file system overview, select the **Replication relationships** tab.
6. In the Replication relationships tab, select the actions menu of the replication relationship to initialize.
7. Select **Initialize**.
8. In the Initialize relationship dialog, select **Initialize**.

Protect your data with NetApp Autonomous Ransomware Protection with AI

Protect your data with NetApp Autonomous Ransomware Protection with AI (ARP/AI), a feature that uses workload analysis in NAS (NFS/SMB) environments to detect and warn about abnormal activity that might be a ransomware attack. When an attack is suspected, ARP/AI also creates new, immutable snapshots from which you can restore your data.

About this task

Use ARP/AI to protect against denial-of-service attacks where the attacker withholds data until a ransom is paid. ARP/AI offers real-time ransomware detection based on:

- Identification of the incoming data as encrypted or plaintext.
- Analytics that detect:
 - **Entropy**: An evaluation of the randomness of data in a file
 - **File extension types**: An extension that does not conform to the normal extension type
 - **File IOPS**: A surge in abnormal volume activity with data encryption

ARP/AI can detect the spread of most ransomware attacks after only a small number of files are encrypted, take action automatically to protect data, and alert you that a suspected attack is happening.

The ARP/AI feature automatically updates according to the ONTAP version that Amazon FSx for NetApp ONTAP runs so you don't have to make manual updates.

Learning and active modes

ARP/AI operates first in *learning mode* and then automatically switches to *active mode*.

- **Learning mode:** When you enable ARP/AI it runs in *learning mode*. In learning mode, the FSx for ONTAP file system develops an alert profile based on the analytic areas: entropy, file extension types, and file IOPS. After the file system runs ARP/AI in learning mode for enough time to assess workload characteristics, Workload Factory automatically switches to ARP/AI to *active mode* and starts protecting your data.
- **Active mode:** After ARP/AI switches to *active mode*, FSx for ONTAP creates ARP/AI snapshots to protect the data if a threat is detected.

In active mode, if a file extension is flagged as abnormal, you should evaluate the alert. You can act on the alert to protect your data or you can mark the alert as a false positive. Marking an alert as a false positive updates the alert profile. For example, if the alert is triggered by a new file extension and you mark the alert as a false positive, you will not receive an alert the next time that file extension is observed.

FlexVol volumes containing a block device start ARP/AI in active mode.

Unsupported configurations

The following configurations don't support the use of ARP/AI.

- iSCSI volumes
- NVMe volumes

Enable ARP/AI for a file system or a volume

Enabling ARP/AI for a file system adds protection for all existing NAS and newly created NAS (NFS/SMB) volumes automatically. You can also enable ARP/AI for individual volumes.

After enabling ARP/AI, if an attack occurs and you identify the attack is real, Workload Factory automatically sets up a snapshot policy that takes up to six snapshots every four hours. Each snapshot is locked for 2-5 days.

Before you begin

To enable ARP/AI for a file system or a 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.

Enable ARP/AI for a file system

Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the actions menu of the file system to enable ARP/AI and then select **Manage**.
5. Under Information, select the pencil icon next to **Autonomous Ransomware Protection**. The pencil icon appears next to the arrow when the mouse hovers over the **Autonomous Ransomware Protection** row.
6. From the NetApp Autonomous Ransomware Protection with AI (ARP/AI) page, do the following:
 - a. Enable or disable the feature.
 - b. **Automatic snapshot creation**: Select the maximum number of snapshots to retain and the interval of time between taking snapshots. The default is 6 snapshots every 4 hours.
 - c. **Immutable snapshots**: Select the default retention period in hours and the maximum number of days to retain immutable snapshots. Enable this option to ensure that snapshots cannot be deleted or modified until the specified retention period ends.
 - d. **Detection**: Optionally, select any of the following parameters to automatically scan and detect anomalies.
7. Accept the statement to proceed.
8. Select **Apply** to save the changes.

Enable ARP/AI for a volume

Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the actions menu of the file system to enable ARP/AI and then select **Manage**.
5. From the Volumes tab, select the actions menu of the volume to enable ARP/AI, then **Data protection actions**, and then **Manage ARP/AI**.
6. In the Manage ARP/AI dialog, do the following:
 - a. Enable or disable the feature.
 - b. **Detection**: Optionally, select any of the following parameters to automatically scan and detect anomalies.
7. Accept the statement to proceed.
8. Select **Apply** to save the changes.

Validate ransomware attacks

Determine if an attack is a false alarm or a genuine ransomware incident.

Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the file system to validate ransomware attacks for.
5. From the file system overview, select the **Volumes** tab.
6. Select **Analyze attacks** from the Autonomous Ransomware Protection tile.
7. Download the attack events report to review if any files or folders were compromised and then decide if an attack has occurred.
8. If no attack occurred, select **False alarm** for the volume in the table and then select **Close**.
9. If an attack has occurred, select **Real attack** for the volume in the table. The Restore compromised volume data dialog opens. You can proceed to [recover your data](#) immediately or select **Close** and come back to complete the recovery process later.

Recover data after a ransomware attack

When an attack is suspected, the system takes a volume snapshot at that point in time and locks that copy. If the attack is confirmed later, the affected files or the entire volume can be restored using the ARP/AI snapshot.

Locked snapshots cannot be deleted until the retention period ends. However, if you decide later to mark the attack as a false positive, the locked copy will be deleted.

With the knowledge of the affected files and the time of attack, it is possible to selectively recover the affected files from various snapshots rather than simply reverting the whole volume to one of the snapshots.

Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the file system to recover data for.
5. From the file system overview, select the **Volumes** tab.
6. Select **Analyze attacks** from the Autonomous Ransomware Protection tile.
7. If an attack has occurred, select **Real attack** for the volume in the table.
8. In the Restore compromised volume data dialog, follow the instructions to restore at the file-level or at the volume-level. In most cases, you'll restore files rather than an entire volume.
9. After you complete the restore, select **Close**.

Result

The compromised data has been restored.

Clone a volume in NetApp Workload Factory

Clone a volume in NetApp Workload Factory to make a read/write volume of the original volume for testing.

The clone reflects the current, point-in-time state of the data. You can also use clones to give additional users access to data without giving them access to production data.

About this task

Volume cloning is only supported for FlexClone volumes.

When a volume is cloned, a writeable volume is created with references to snapshots from the parent volume. Clone creation occurs in seconds. The cloned data doesn't reside on the volume clone but instead resides on the parent volume. Any new data written to the volume after clone creation resides on the clone.

For a cloned volume to contain all data from the parent volume and any new data added to the clone after creation, you'll need to [split the clone](#) from the parent volume. Additionally, you can't delete a parent volume if it has a clone. A clone must be split before a parent volume can be deleted.

Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the actions menu of the FSx for ONTAP file system which contains the volume to clone then select **Manage**.
5. From the Overview tab of the file system, select the **Volumes** tab.
6. In the Volumes tab, select the actions menu of the volume to clone.
7. Select **Data protection actions**, then **Clone volume**.
8. In the Clone volume dialog, enter a name for the volume clone.
9. Select **Clone**.

Use on-premises ONTAP cluster data in NetApp Workload Factory

Discover and replicate on-premises ONTAP data in NetApp Workload Factory so it can be used to enrich AI knowledge bases.

About this task

To use data from an on-premises ONTAP cluster, you'll first need to discover the on-premises ONTAP cluster. After you've discovered an on-premises ONTAP cluster, you can use the data for any of the following use cases.

Use cases

Note that the primary use case for the GenAI workload is the focus of this series of tasks.

- **GenAI workload:** Replicate on-premises-ONTAP volume data to an FSx for ONTAP file system so that the data can be used to [enrich AI knowledge bases](#).
- **Backup and migration to cloud:** Replicated on-premises ONTAP volume data to an FSx for ONTAP file system can be used as a backup in the cloud.
- **Data tiering:** After replication, infrequently accessed on-premises ONTAP volume data can be tiered from the SSD storage tier to the capacity pool storage tier.

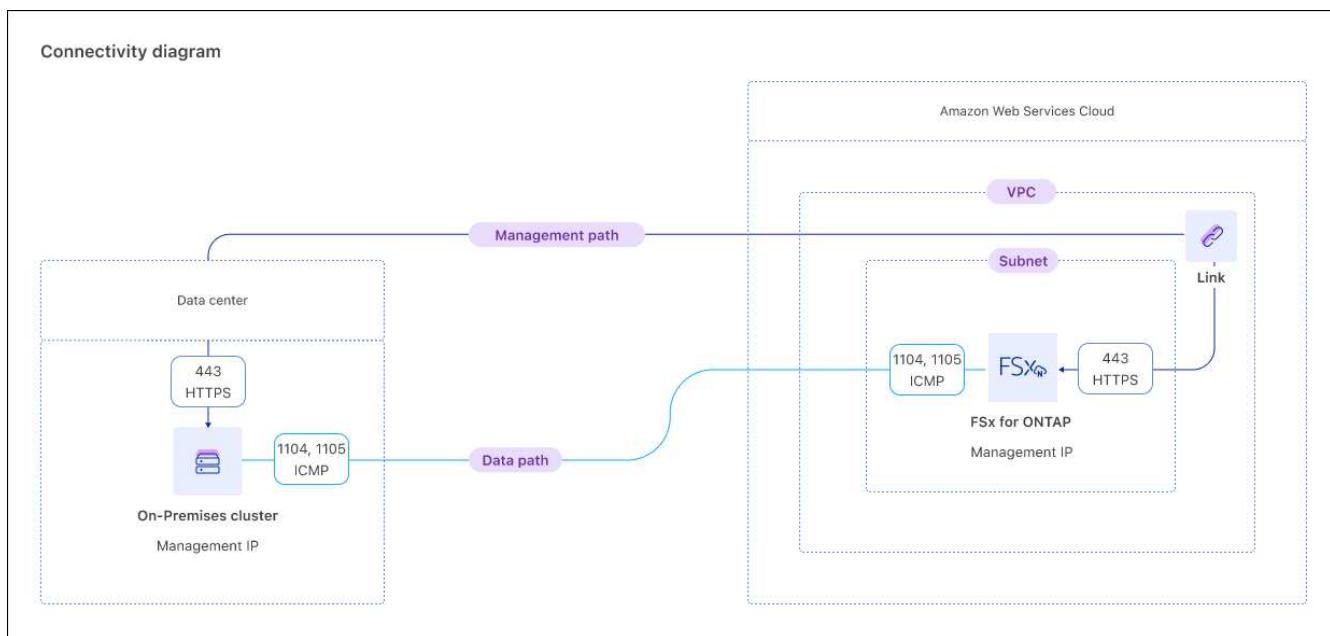
Discover an on-premises ONTAP cluster

Discover an on-premises ONTAP cluster in NetApp Workload Factory so that you can replicate the data to an Amazon FSx for NetApp ONTAP file system.

Before you begin

Make sure you have the following before you begin:

- An FSx for ONTAP file system for replication.
- A connected link to associate with the discovered on-premises cluster. If you don't have a link, you'll need to [create one](#).
- ONTAP user credentials with required permissions.
- On-premises ONTAP version 9.8 and above.
- Connectivity as shown in the following diagram.



Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. Select the **On-premises ONTAP** tab.
4. Select **Discover**.
5. Review the prerequisites and select **Next**.
6. On the Discover ONTAP on-premises page, provide the following under **Cluster configuration**:
 - a. **Link:** Select a link. The link will be associated with the on-premises cluster to create connectivity between the cluster and Workload Factory.
If you haven't created a link, follow the instructions and then return to this operation and select the link.
 - b. **Cluster IP address:** Provide the IP address for the on-premises ONTAP cluster to replicate.
 - c. **ONTAP credentials:** Enter the ONTAP credentials for the on-premises ONTAP cluster. Make sure the

user has the required permissions.

7. Select **Discover** to start the discovery process.

Result

The on-premises ONTAP cluster is discovered and now appears in the **On-Premises ONTAP** tab.

You can now view the data in your on-premises ONTAP cluster and [replicate the data to an FSx for ONTAP file system](#).

Replicate volume data from an on-premises ONTAP cluster

Replicate volume data from an on-premises ONTAP cluster to an FSx for ONTAP file system. After replication, the data can be used to enrich AI knowledge bases.

Before you begin

- You must discover an on-premises ONTAP cluster to replicate its volume data.
- You must have an available FSx for ONTAP file system to be the target for the replication.
- Both the on-premises ONTAP cluster and the FSx for ONTAP file system you use for the replication relationship must have an associated 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 using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **On-premises ONTAP**.
4. To find volumes by storage VM, you can **Select storage VM** from the dropdown.
5. Select one or more volumes to replicate and then select **Replicate**.
6. On the Create replication page, under Replication target, provide the following:
 - a. **FSx for ONTAP file system**: Select credentials, region, and FSx for ONTAP file system name for the target FSx for ONTAP file system.
 - b. **Storage VM name**: Select the storage VM from the dropdown menu.
 - c. **Volume name**: The target volume name is generated automatically with the following format {OriginalVolumeName}_copy. You can use the auto-generated volume name or enter another volume name.
 - d. **Tiering data**: Select the tiering policy for the data stored in the target volume.
 - **Auto**: The default tiering policy when creating a volume using the Workload Factory FSx for ONTAP user interface. Tiers all cold data that includes user data and snapshots to the capacity pool storage tier for a specific time period.
 - **Snapshot only**: Tiers only snapshot data to the capacity pool storage tier.
 - **None**: Keeps all your volume's data on the primary storage tier.
 - **All**: Marks all user data and snapshot data as cold and stores it in 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.

For more information about volume tiering policies, refer to [Volume storage capacity](#) in AWS FSx for NetApp ONTAP documentation.

e. **Max transfer rate:** Select **Limited** and enter the max transfer limit in MiB/s. Alternatively, select **Unlimited**.

Without a limit, network and application performance might decline. Alternatively, we recommend an unlimited transfer rate for FSx for ONTAP file systems for critical workloads, for example, those that are used primarily for disaster recovery.

7. Under Replication settings, provide the following:

a. **Replication interval:** Select the frequency that snapshots are transferred from the source volume to the target volume.

b. **Long-term retention:** Optionally, enable snapshots for long-term retention.

If you enable long-term retention, then select an existing policy or create a new policy to define the snapshots to replicate and the number to retain.

- For an existing policy, select **Choose an existing policy** and then select the existing policy from the dropdown menu.
- For a new policy, select **Create a new policy** and provide the following:
 - **Policy name:** Enter a policy name.
 - **Snapshot policies:** In the table, select the snapshot policy frequency and the number of copies to retain. You can select more than one snapshot policy.

8. Select **Create**.

Result

The replication relationship appears in the **Replication relationships** tab in the target FSx for ONTAP file system.

Remove an on-premises ONTAP cluster from NetApp Workload Factory

Remove an on-premises ONTAP cluster from NetApp Workload Factory when needed.

Before you begin

You must [delete all existing replication relationships](#) for any volumes in the on-premises ONTAP cluster before removing the cluster so that no broken relationships remain.

Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **On-premises ONTAP**.
4. Select the on-premises ONTAP cluster to remove.
5. Select the actions menu and select **Remove from Workload Factory**.

Result

The on-premises ONTAP cluster is removed from NetApp Workload Factory.

Protect your data with a cyber vault

A cyber vault volume is an isolated, secure storage location used to store backup copies of your data, protecting it from ransomware attacks and other cyber threats. As part of vault creation, you'll create a cyber vault volume, disable all client protocols, and set up a replication relationship between the source volume and the cyber vault volume, and create immutable snapshots on the cyber vault volume.

What is a cyber vault?

A cyber vault is a specific data protection technique that involves storing critical data in an isolated environment, separate from the primary IT infrastructure.

The cyber vault is an "air-gapped", immutable, and indelible data repository that is immune to threats affecting the main network, such as malware, ransomware, or even insider threats. A cyber vault can be achieved with immutable and indelible snapshots.

Air-gapping backups that use traditional methods involve creating space and physically separating the primary and secondary media. By moving the media offsite and/or severing connectivity, bad actors have no access to the data. This protects the data but can lead to slower recovery times.

FSx for ONTAP cyber vaults

Amazon FSx for NetApp ONTAP is supported as a cyber vault source and target.

Implementation

Workload Factory provides assistance in creating a cyber vault architecture. After you contact NetApp to express your interest in implementing a cyber vault, a NetApp specialist contacts you to discuss your requirements.

Send an email to ng-FSx-CyberVault@netapp.com to get started.

Related information

For more information about cyber vaults and how to set up this architecture, refer to the [ONTAP cyber vault documentation](#).

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—with 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.