



# **Amazon FSx for NetApp ONTAP for NetApp Workload Factory documentation**

Amazon FSx for NetApp ONTAP

NetApp  
April 13, 2026

# Table of Contents

- Amazon FSx for NetApp ONTAP for NetApp Workload Factory documentation . . . . . 1
- Release notes . . . . . 2
  - What’s new with Amazon FSx for NetApp ONTAP . . . . . 2
    - 06 April 2026 . . . . . 2
    - 01 March 2026 . . . . . 2
    - 16 February 2026 . . . . . 4
    - 01 February 2026 . . . . . 4
    - 04 January 2026 . . . . . 5
    - 04 December 2025 . . . . . 6
    - 27 November 2025 . . . . . 7
    - 02 November 2025 . . . . . 7
    - 06 October 2025 . . . . . 9
    - 05 October 2025 . . . . . 9
    - 09 September 2025 . . . . . 10
    - 03 August 2025 . . . . . 12
    - 29 June 2025 . . . . . 14
    - 08 June 2025 . . . . . 15
    - 03 June 2025 . . . . . 16
    - 04 May 2025 . . . . . 16
    - 30 March 2025 . . . . . 18
    - 02 March 2025 . . . . . 19
    - 02 February 2025 . . . . . 20
    - 05 January 2025 . . . . . 21
    - 1 December 2024 . . . . . 21
    - 3 November 2024 . . . . . 22
    - 29 September 2024 . . . . . 23
    - 1 September 2024 . . . . . 23
    - 4 August 2024 . . . . . 23
    - 7 July 2024 . . . . . 23
  - Known limitations of Amazon FSx for NetApp ONTAP in NetApp Workload Factory . . . . . 24
    - Region support . . . . . 24
    - Adding HA pairs limitations . . . . . 24
    - Throughput capacity region support . . . . . 24
    - Capacity management . . . . . 24
    - Storage VMs . . . . . 25
    - iSCSI protocol support . . . . . 25
    - Data protection . . . . . 25
    - Storage savings calculator . . . . . 25
    - AWS Secrets Manager support . . . . . 25
    - Amazon S3 access points limitation . . . . . 25
- Get started . . . . . 27
  - Learn about Amazon FSx for NetApp ONTAP in NetApp Workload Factory . . . . . 27
  - Features . . . . . 27

Additional features in Workload Factory	28
Tools to use NetApp Workload Factory	28
Cost	28
Regions	28
Getting help	29
Quick start for Amazon FSx for NetApp ONTAP in NetApp Workload Factory	29
Create an FSx for ONTAP file system in NetApp Workload Factory	30
Create an FSx for ONTAP file system	30
Security group details	35
Use Amazon FSx for NetApp ONTAP	38
Explore savings with FSx for ONTAP in NetApp Workload Factory	38
Calculator options	38
Explore savings via customization	38
Explore savings for detected storage environments	40
Deploy FSx for ONTAP file systems	42
Track costs for your resources in NetApp Workload Factory	42
Use links	43
Learn about NetApp Workload Factory links	43
Connect to an FSx for ONTAP file system with a Lambda link	44
Manage Workload Factory links	51
Discover cache volumes in Workload Factory	54
Manage volumes	55
Create an FSx for ONTAP volume in Workload Factory	55
Access your FSx for ONTAP file system data	60
Create block storage resources	61
Create an initiator group for a file system in NetApp Workload Factory	61
Create a block device for a file system in NetApp Workload Factory	62
Create a storage VM for an FSx for ONTAP file system	64
Create a storage VM	64
Protect your data	65
Types of data protection in NetApp Workload Factory	65
Use snapshots	67
Use backups to object storage	70
Use replication	72
Protect your data with NetApp Autonomous Ransomware Protection with AI	77
Clone a volume in NetApp Workload Factory	80
Use on-premises ONTAP cluster data in NetApp Workload Factory	81
Protect your data with a cyber vault	84
Administer and monitor	86
Monitor storage operations with Tracker in NetApp Workload Factory	86
Track and monitor operations	86
View API request	87
Retry a failed operation	87
Edit and retry a failed operation	87
Implement file system best practices	88

Configuration analysis for FSx for ONTAP file systems . . . . .	88
Implement well-architected file system configurations . . . . .	91
Analyze FSx for ONTAP EMS events in NetApp Workload Factory . . . . .	95
About this task . . . . .	95
Before you begin . . . . .	96
View and analyze EMS events for FSx for ONTAP . . . . .	97
Volume administration . . . . .	97
Enable volume autogrow in Workload Factory . . . . .	97
Adjust volume capacity in NetApp Workload Factory . . . . .	98
Check and rebalance volume capacity . . . . .	99
Manage immutable files for a volume in NetApp Workload Factory . . . . .	102
Manage volume tags in NetApp Workload Factory . . . . .	103
Manage FSx for ONTAP cache volumes with NetApp Workload Factory . . . . .	104
Change the tiering policy of a volume in NetApp Workload Factory . . . . .	106
Update storage efficiency setting of a volume . . . . .	107
Manage the NFS export policy for a volume in NetApp Workload Factory . . . . .	108
Manage the SMB/CIFS share for a volume in Workload Factory . . . . .	109
Manage S3 access points . . . . .	111
Split a cloned volume in NetApp Workload Factory . . . . .	124
Delete a volume in NetApp Workload Factory . . . . .	125
Block storage administration . . . . .	125
Manage block storage for a file system in NetApp Workload Factory . . . . .	125
File system administration . . . . .	129
Manually adjust file system capacity in Workload Factory . . . . .	129
Manage file system capacity and inodes automatically in Workload Factory . . . . .	130
Manage FSx for ONTAP file system tags in NetApp Workload Factory . . . . .	133
Reset the fsxadmin password in NetApp Workload Factory . . . . .	133
Delete a file system in NetApp Workload Factory . . . . .	134
Storage VM administration . . . . .	134
Replicate a storage VM to another FSx for ONTAP file system . . . . .	134
Configure and update Active Directory for a storage VM . . . . .	136
Manage storage VM tags in NetApp Workload Factory . . . . .	137
Reset the storage VM password in NetApp Workload Factory . . . . .	137
Delete a storage VM in NetApp Workload Factory . . . . .	137
Data protection administration . . . . .	138
Snapshots . . . . .	138
Backups . . . . .	144
Replication . . . . .	144
Performance administration . . . . .	151
Provision SSD IOPS for an FSx for ONTAP file system . . . . .	151
Update throughput capacity for a file system . . . . .	152
Reference . . . . .	153
Performance for FSx for ONTAP in NetApp Workload Factory . . . . .	153
Security for FSx for ONTAP in NetApp Workload Factory . . . . .	153
Knowledge and support . . . . .	154

Register for support .....	154
Support registration overview .....	154
Register your account for NetApp support .....	154
Get help for FSx for ONTAP for Workload Factory .....	156
Get support for FSx for ONTAP .....	156
Use self-support options .....	156
Create a case with NetApp support .....	156
Manage your support cases (Preview) .....	159
Legal notices for NetApp Workload Factory .....	162
Copyright .....	162
Trademarks .....	162
Patents .....	162
Privacy policy .....	162
Open source .....	162

# Amazon FSx for NetApp ONTAP for NetApp Workload Factory documentation

# Release notes

## What's new with Amazon FSx for NetApp ONTAP

Learn what's new with Amazon FSx for NetApp ONTAP.

### 06 April 2026

#### Well-architected analysis updates

Workload Factory analyzes your FSx for ONTAP file systems for **inactive NAS volumes**. The configuration analysis identifies NAS volumes that are not actively used and recommends actions to optimize storage utilization and reduce costs.

[Implement well-architected FSx for ONTAP file systems](#)

#### Journal table feature for S3 access points

The Journal table feature provides a way to manage FSx ONTAP S3 object storage at scale, allowing you to easily audit user access events and operations on objects in your buckets. With Workload Factory, you can set up automatic scans of your buckets and S3 tables to keep a record of events, showing all objects, their details, and user actions.

[Set up the journal table infrastructure for NetApp Workload Factory](#)

#### Support for managing existing igroups

Workload Factory supports managing existing igroups. You can add or remove block devices and host initiators to deliver your deployed workload on time.

[Manage block storage in NetApp Workload Factory](#)

#### Support for specifying the aggregate for FlexVol volumes

Workload Factory supports creating FlexVol volumes on a chosen aggregate. This allows you to split the capacity and performance load across all the file system resources.

[Create a FlexVol volume in NetApp Workload Factory](#)

#### Last access time for block devices

Workload Factory reports the last access time for block devices. The time is based on when you last accessed the igroup in Workload Factory. Knowing when each block device was last used allows you to optimize your resources.

### 01 March 2026

#### Well-architected analysis updates

Workload Factory analyzes your FSx for ONTAP file systems for the following configurations:

- Underutilized SSD capacity: checks whether the SSD capacity tier is underutilized and recommends

decreasing SSD capacity to optimize costs.

- Schedule volume backups: Workload Factory improves the configuration recommendation and adds options to turn on FSx for ONTAP backup or AWS Backup. FSx for ONTAP backup applies a default policy to all volumes in a file system, creating daily backups kept for up to 90 days. AWS Backup lets you choose different retention periods for each volume in the backup plan.

## **Workload Factory supports AI & ML use case for replication**

Workload Factory supports replicating on-premises ONTAP data to FSx for ONTAP and creating S3 access points for AWS Artificial Intelligence (AI) and machine learning (ML) services to accelerate AI/ML workflows.

[Replicate on-premises data for AI & ML use cases](#)

## **Editing existing S3 access points**

You can change the user and user type for an existing S3 access point attached to a volume. You can also enable or disable metadata for the access point.

[Edit S3 access points in NetApp Workload Factory](#)

## **Added block storage diagram**

If you're using block storage, you can view the block storage diagram to visualize how block devices connect to nodes and igroups. The diagram helps you understand and manage your block storage.

[Manage block storage in NetApp Workload Factory](#)

## **Update to automatic capacity management feature**

This update separates the feature into two distinct features: automatic capacity management and capacity notifications.

Automatic capacity management automatically adds or removes file system capacity to maintain optimal performance. New options include incremental and adaptive modes. In incremental mode, the system increases capacity by a fixed amount at thresholds; in adaptive mode, it uses historical data to predict and adjust capacity.

Capacity notifications alert you whenever a capacity adjustment occurs and when file systems reach set limits, so you can manage storage before issues arise.

[Learn about automatic capacity management and capacity notifications in NetApp Workload Factory](#)

## **New Ask Me bookmark available in Workload Factory**

We added a new Ask Me bookmark on every screen in Workload Factory console. This improvement makes it easier and faster for you to access Ask Me whenever you need help. Ask Me is our AI assistant. You can ask questions about your workload environments, get personalized insights, and review previous conversations.

You can open Ask Me from any page by clicking the new bookmark. It launches Ask Me in a side panel without interrupting your current work and offers quick explanations and recommendations related to what you're currently doing.

[Learn more about Ask Me](#)

## 16 February 2026

### Support for storage VM migration

NetApp Workload Factory now supports migration for storage VMs. This feature allows migration of ONTAP storage system data and configurations from on-premises ONTAP systems or first-generation FSx for ONTAP file systems to second-generation FSx for ONTAP file systems. You can replicate storage VM data and configuration settings to move to the new file system with minimal downtime and disruption to users and applications.

To use this feature, [create a replication relationship](#) and select **Migration** as the use case. To complete the migration process, you must [cut over the storage VM and its replicated volumes](#) immediately to permanently migrate data and storage VM configuration settings to the target FSx for ONTAP file system.

## 01 February 2026

### Home page includes well-architected issues and EMS events for Storage

The NetApp Workload Factory home page includes a Focus tile where well-architected issues and FSx for ONTAP Emergency Management System (EMS) events appear for your workloads. From there, you can navigate to the Storage workload to view the well-architected status or the events of all FSx for ONTAP file systems in your storage environment.

### Support for on-premises data replication using an S3 access point

Workload Factory supports replicating on-premises ONTAP data to the cloud for integration with AWS GenAI, ML, and analytics. You can replicate your on-premises data to an NFS or SMB/CIFS volume using an S3 access point.

[Replicate on-premises data using an S3 access point](#)

### S3 access point enhancements in Storage

Several enhancements have been made to the S3 access point management capabilities in the Storage workload for NetApp Workload Factory. You can input network configuration details for your S3 access points and add S3 access point tags. Additional enhancements include the ability to view S3 bucket details and perform more actions for managing S3 access points.

### S3 bucket details available in Storage

The new Inventory table automatically scans your AWS S3 buckets and populates S3 tables to give you a clear snapshot of all objects, their metadata, attributes, and tags. Access to these details helps you maintain control, visibility, and trust in the data you're responsible for, while reducing operational overhead. You can turn on the feature when creating and attaching S3 access points or when editing existing access points.

[Create and attach an S3 access point for a volume in NetApp Workload Factory](#)

### Additional management operations for S3 access points

NetApp Workload Factory provides additional management operations for S3 access points. You can view access point details, modify existing S3 access points, and add or remove S3 access point tags from the NetApp Workload Factory interface, streamlining your object storage management tasks.

## Well-architected analysis updates

Workload Factory analyzes your FSx for ONTAP file systems for the following configurations:

- Optimize cache volume size: checks whether volume autosize and scrubbing are enabled on cache volumes to maintain optimal size and focus the cache on hot data for peak efficiency.
- Inactive block devices: recommends archiving block device data or deleting a block device if it hasn't been used for seven consecutive days.
- Storage VM logical reporting: checks whether the default reporting setting for a storage VM is set to logical which provides better visibility into storage usage at the volume level.

## Additional cards for block devices

Three new cards have been added to the Block devices tab in the Storage workload to provide quick insights into block device usage and protection status:

- Storage efficiencies: displays used and available storage capacity; used capacity is broken down by SSD and capacity pool storage tiers.
- Protected devices: displays the percentage of block devices with snapshots, remote replication, NetApp Autonomous Ransomware Protection (ARP/AI), and backups.
- Inactive block devices: displays if any block devices haven't been used for seven consecutive days, helping you identify and manage unused resources effectively. From here, you can [reclaim space for unused block devices](#).

## Support for initiator group creation and management

NetApp Workload Factory supports creating and managing initiator groups (igroups) for block storage in FSx for ONTAP file systems. Initiator groups connect block devices (LUNs) to the compute resources that are allowed to access them, providing a permission layer for block storage in SAN environments.

- [Create an initiator group in NetApp Workload Factory](#)
- [Manage existing initiator groups in NetApp Workload Factory](#)

## 04 January 2026

### Well-architected analysis updates

Workload Factory analyzes your FSx for ONTAP file systems for the following configurations:

- NetApp Autonomous Ransomware Protection (ARP/AI) disabled includes block devices: checks whether ARP/AI is disabled on block device volumes
- Cache relationship write mode: checks whether the write mode is optimal for the cache volume workload
- Unnecessary backup deletion: checks whether backups are outdated or unnecessary that can be deleted to reduce costs

[View the well-architected status of your FSx for ONTAP file systems](#)

### Ask me AI assistant home page integration

The Workload Factory console home page embeds the Ask me AI assistant, enabling you to ask questions about your own storage estate, get personalized insights directly from your environment, and refer to previous

conversations. You can interact with Ask me to understand your workloads, troubleshoot issues, and learn more about Workload Factory — all without leaving the console.

### **Use of IAM user principal in Lambda link resource-based permission policies**

Lambda links that are used to connect between your Workload Factory account and one or more FSx for ONTAP file systems to perform advanced ONTAP operations, now use the IAM user principal for resource-based policy permissions. This change provides better alignment with industry best practices for AWS resource access.

### **Analysis screen added for the AI analyzer for EMS events**

A new *Analysis* screen has been added to the Storage menu. From this screen, you can use the AI analyzer for FSx for ONTAP EMS events feature.

### **Block device enhancements in NetApp Workload Factory**

The following enhancements have been made for block devices.

#### **Block device creation**

NetApp Workload Factory supports creating block devices using the iSCSI protocol on FSx for ONTAP file systems so that you can better support your line of business (LOB) applications from the Workload Factory console.

#### **Block device management enhancements**

NetApp Workload Factory includes the following enhancements for [managing block devices](#). You can now perform the following tasks from the Workload Factory console:

- Manage client access
- Archive block device data
- Delete a block device

#### **Support for ARP/AI on FlexVol volumes containing block devices**

You can enable [NetApp Autonomous Ransomware Protection with AI \(ARP/AI\)](#) on FlexVol volumes that contain block devices. Enabling ARP/AI detects ransomware attacks using AI and aids in data recovery.

## **04 December 2025**

### **Support for AWS S3 access points for FSx for ONTAP**

NetApp Workload Factory supports AWS S3 access points for your FSx for ONTAP file systems. You can create volumes using S3 access points, assign S3 access points to an existing volume, and manage S3 access points from the Workload Factory console. Using an S3 access point, you can access file data residing on SMB/CIFS or NFS volumes via the AWS S3 APIs. This allows you to integrate your existing data with GenAI, ML, and analytics from AWS services that support S3 access points.

- [Create a volume using S3 access points](#)
- [Manage S3 access points for a volume](#)

## 27 November 2025

### Block device support in NetApp Workload Factory

Manage your block devices more effectively with the newly introduced block device support in NetApp Workload Factory. This feature allows you to view details and increase capacity for iSCSI LUNs, providing enhanced flexibility for your storage needs.

[Manage block devices in Workload Factory](#)

### Well-architected analysis updates

Workload factory analyzes your FSx for ONTAP file systems for the following configurations:

- Unnecessary snapshot deletion: checks whether volumes have outdated and unnecessary snapshots that can be deleted to reduce costs.
- FlexGroup volumes rebalance: checks whether FlexGroup volumes are evenly balanced across their member volumes to ensure optimal performance.

[View the well-architected status of your FSx for ONTAP file systems](#)

### AI analyzer for EMS events in NetApp Workload Factory

NetApp Workload Factory introduces an AI-powered analyzer for ONTAP Event Management System (EMS) events. This feature helps you quickly identify and troubleshoot issues by providing insights and recommendations based on the analysis of EMS event data.

[Analyze EMS events in Workload Factory](#)

### Monitor cost and usage trends for FSx for ONTAP file systems

You can monitor cost and usage trends for your FSx for ONTAP file systems directly from the NetApp Workload Factory console. This feature provides storage consumption and cost metrics as well as itemized costs, helping you optimize your resource allocation and budget planning.

[Track costs for FSx for ONTAP file systems in Workload Factory](#)

### Manage FSx tags for a file system in NetApp Workload Factory

Easily manage your FSx tags for a file system directly from the NetApp Workload Factory console. This feature allows you to add, edit, or remove tags, enabling better organization and categorization of your FSx for ONTAP file systems.

[Manage FSx tags in Workload Factory](#)

### Adjust cache capacity for FSx for ONTAP file systems

You can increase and decrease capacity for cache volumes from the Workload Factory console.

[Manage cache volumes in Workload Factory](#)

## 02 November 2025

## Cache volume management

You can perform the following cache volume management operations from within the Workload Factory console:

- Edit the cache name
- Increase the capacity of a cache volume
- Edit the mount path or export policy for a cache volume
- Change the caching method, or mode, for a cache volume
- Prepopulate a cache volume
- Delete a cache volume

### [Manage cache volumes](#)

## Automatic inode management available

You can enable automatic inode management without the need to enable automatic capacity management.

### [Enable automatic inode management](#)

## Threshold warning setting for capacity and inode usage

Threshold warnings are available for both capacity and inode usage. You can set these thresholds to when enabling automatic capacity or inode management. To use this setting, you'll need to configure notifications using the [NetApp Workload Factory notification service](#).

## Volume size decrease available

You can decrease the size of NFS and SMB/CIFS volumes in NetApp Workload Factory. This feature allows for better management of storage resources by enabling you to reduce the size of volumes that are no longer needed at their current capacity.

### [Decrease the capacity of a volume](#)

## Enhanced FSx for ONTAP resource state

Workload Factory has enhanced the "misconfigured" resource state to include an explanation of the actual issue for the resource.

## Well-architected analysis updates

Workload factory analyzes your FSx for ONTAP file systems for the following configurations:

- Volume utilization nearing full: checks whether any volumes are using 80% or more of their file capacity. This helps you identify volumes that may need additional capacity.
- Unauthorized access to volumes: checks whether an iSCSI volume is accessible using an NFS or SMB/CIFS mount path and allows you to remove unauthorized access to the volume to avoid security risks.

## Permissions changes for Workload Factory for Storage

Workload Factory provides more clarity about the permissions it requires for specific actions and granularity for

selecting only the permissions you need. When you add credentials, you'll have three permissions options to choose from instead of the previous permissions model which was *read-only* and *read/write*. The new permissions model breaks up the permission policies as follows:

- *View, planning, and analysis*: View FSx for ONTAP file systems, learn about system health, get the well-architected analysis for your systems, and explore savings.
- *Operations and remediation*: Perform operational tasks like adjust file system capacity and fix issues for your file system configurations.
- *File system creation and deletion*: Create and delete FSx for ONTAP file systems and storage VMs.

When adding credentials, you can select one or more of these permission policies based on the level of access you want to provide to Workload Factory.

[Workload Factory permissions reference](#)

### **FSx for ONTAP cyber vault support**

You can create a cyber vault using FSx for ONTAP as a source or target in the cyber vault architecture. Cyber vaults provide a secure and isolated environment for storing critical data, protecting it from ransomware and other cyber threats.

[Set up a cyber vault with FSx for ONTAP](#)

## **06 October 2025**

### **BlueXP workload factory now NetApp Workload Factory**

BlueXP has been renamed and redesigned to better reflect the role it has in managing your data infrastructure. As a result, BlueXP workload factory has been renamed to NetApp Workload Factory.

## **05 October 2025**

### **Optimize savings in the Storage calculator for Amazon Elastic Block Store (EBS)**

Workload Factory can analyze your EBS performance usage and then suggest the best and most cost-efficient FSx for ONTAP configuration so that you can save more by switching to FSx for ONTAP.

[Explore savings for detected storage environments in the Workload Factory console](#)

### **Quick access to resource screen from file system inventory**

You can quickly navigate to an FSx for ONTAP file system resource screen by selecting the file system name, now a hyperlink, from the FSx for ONTAP inventory.

### **Discover cache relationships in the Workload Factory console**

If you have *cache* relationships between FSx for ONTAP file system and another type of ONTAP storage (on-premises system, Cloud Volumes ONTAP, and FSx for ONTAP), you can discover and view them from the Workload Factory console. This allows you to better understand data flows, optimize cache utilization, and improve efficiency across distributed environments.

[Discover and view cache relationships in the Workload Factory console](#)

## Well-architected analysis update

Workload factory now analyzes your FSx for ONTAP file systems for the following configuration:

Volume file capacity utilization threshold: checks whether the file capacity thresholds are set to 80% or lower. This helps you avoid running out of space on your file systems.

[View the well-architected status of your FSx for ONTAP file systems](#)

## Improvements to actions for configuration issues

From the **Well-architected analysis** tab in the dashboard for an FSx for ONTAP file system, instead of dismissing an entire configuration for a file system, you can also select one or more volumes within a file system to fix, dismiss, or reactivate.

## Additional notification for Storage

The NetApp Workload Factory notification service includes the notification for well-architected configuration issues on a weekly basis.

[Notification types and messages in the Workload Factory setup and administration documentation](#)

## Immutable files support privileged delete

With this feature, you can configure privileged delete access for immutable files in your FSx for ONTAP file systems. This allows you to protect critical data from accidental or malicious deletion while still enabling authorized users to override the lock and delete these files as needed. Enabling privileged delete is available during volume creation or for existing volumes.

## 09 September 2025


### Storage inventory report enhancements

Workload factory has enhanced the data reported for your FSx for ONTAP file systems. The downloadable report from the FSx for ONTAP inventory page includes the following new columns:

- SSD used: shows the value of SSD capacity used
- SSD utilization: shows the percentage of SSD capacity in use
- Throughput utilization: shows average and peak utilization for the last 30 days
- IO utilization: shows average and peak IO utilization for the last 30 days
- CPU utilization: shows average and peak CPU utilization for the last 30 days

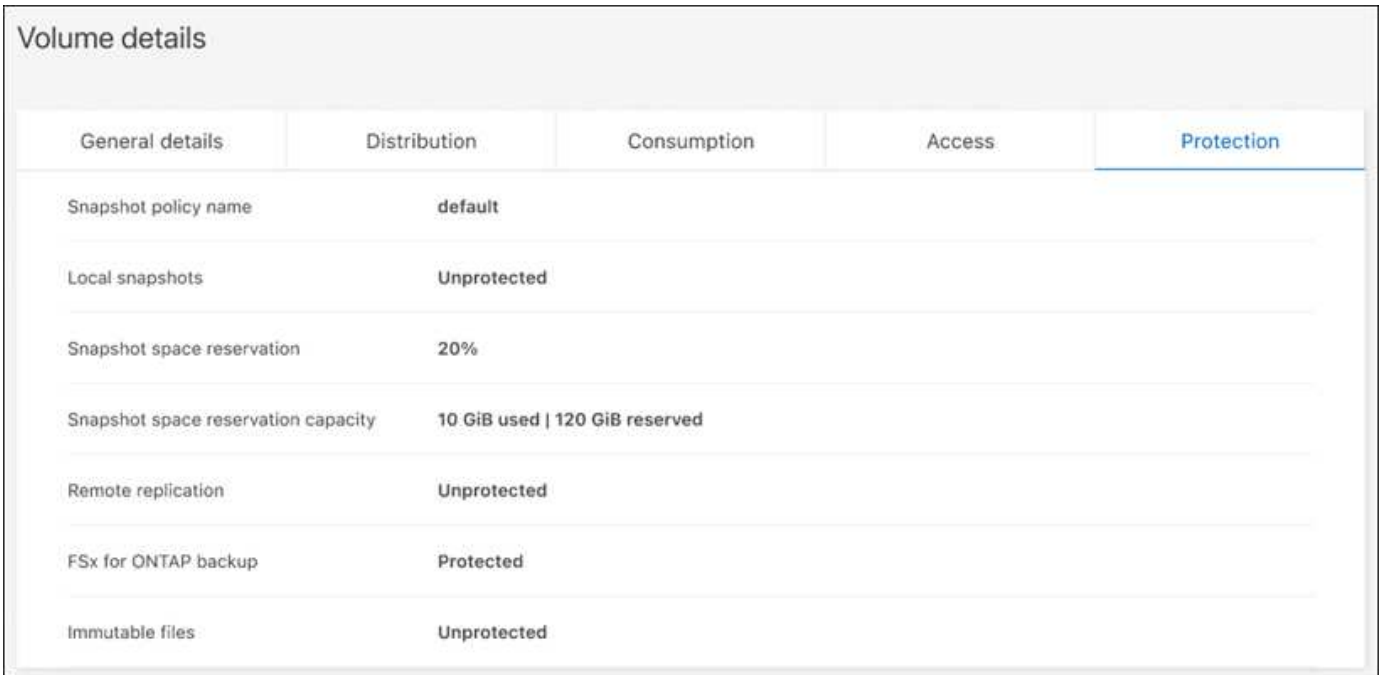
### Snapshot management enhancements

Workload factory has made several enhancements to view volume snapshot details and manage volume snapshots. These enhancements make it easier for you to understand the status of your snapshots and protect your data.



Additional items appear in Volume details under the Protection tab:

- Snapshot policy name
- Snapshot space reservation
- Snapshot space reservation capacity



The new snapshot management screen is accessible from a volume; it provides information about the snapshot policy for the volume and includes a table with all volume snapshots. The table displays the following snapshot details: creation time, size, expiry time, immutable snapshot protection, and labels. From the management screen, you can change the snapshot policy for the volume, create a snapshot manually, and edit, access, restore, and delete snapshots.



- Volume-level ARP/AI: You can now enable ARP/AI at the volume level, allowing you to protect specific volumes within your FSx for ONTAP file systems.
- Automatic snapshot creation: You can set the ARP/AI policy to take automatic snapshots and define how often snapshots are taken for volumes with ARP/AI enabled, enhancing your data protection strategy.
- Immutable snapshots: ARP/AI now supports immutable snapshots, which cannot be deleted or modified, providing an additional layer of security against ransomware attacks.
- Detection: includes various detection methods such as high entropy data rate at volume level, file create rate, file rename rate, file delete rate, and behavioral analysis, and never seen before file extension that help to detect anomalies and potential ransomware attacks.

[Protect your data with NetApp Autonomous Ransomware Protection with AI \(ARP/AI\)](#)

### **Well-architected analysis updates**

Workload factory now analyzes your FSx for ONTAP file systems for the following configurations:

- Long-term retention data reliability: checks whether labels assigned to the snapshot policy of the source volume are identical to the labels assigned to the long-term retention policy. When labels are identical, data replication is reliable between source and target volumes.
- NetApp Autonomous Ransomware Protection with AI (ARP/AI): checks whether ARP/AI is enabled on your file systems. This feature helps you detect and recover from ransomware attacks.

[View the well-architected status of your FSx for ONTAP file systems](#)

### **Dismiss a configuration from the well-architected analysis**

You can now dismiss one or more configurations from the well-architected analysis. This allows you to ignore specific configurations that you don't want to address at the moment.

[Dismiss a configuration from the well-architected analysis](#)

### **Terraform support for link creation**

You can now use Terraform from the Codebox to create a link for association with an FSx for ONTAP file system. This functionality is for users who create links manually.

[Connect to an FSx for ONTAP file system with a Lambda link](#)

### **New region support for exploring savings in Storage**

The following new regions are now supported for exploring savings for Amazon Elastic Block Store (EBS), FSx for Windows File Server, and Elastic File Systems (EFS):

- Mexico
- Thailand

### **Enhancements to SMB/CIFS shares creation and management**

You can now create SMB/CIFS shares which point to directories within a volume. Within the volume, you'll be able to see which shares exist, where the shares are the pointing to, and the permissions granted to specific users and groups.

For data protection volumes, the flow of creating an SMB/CIFS share now includes the creation of a junction path to the volume for mounting purposes.

[Create a CIFS share for a volume](#)

## 29 June 2025

### **BlueXP workload factory notification service support**

The BlueXP workload factory notification service enables workload factory to send notifications to the BlueXP alerts service or to an Amazon SNS topic. Notifications sent to BlueXP alerts appear in the BlueXP alerts panel. When workload factory publishes notifications to an Amazon SNS topic, subscribers to the topic (such as people or other applications) receive the notifications at the endpoints configured for the topic (such as email or SMS messages).

[Configure BlueXP workload factory notifications](#)

### **Storage dashboard enhancements**

The Storage dashboard in the Workload Factory console includes new cards for savings opportunities. The card at the top of the dashboard displays the number of savings opportunities for storage environments running on Amazon Elastic Block Store (EBS), Amazon FSx for Windows File Server, and Amazon Elastic File Systems (EFS). At the bottom of the dashboard, three new cards display savings opportunities by Amazon storage service - EBS, FSx for Windows File Server, and EFS. From all cards, you can explore the savings opportunities in more detail.

From the FSx for ONTAP protection coverage card and replication relationship health card, you can investigate if there are any partially protected volumes in your FSx for ONTAP file systems as well as investigate issues with replication relationships. In both cases, you can take action to resolve the issues.

### **Volume tab enhancements**

The Volumes tab in the Workload Factory console has been enhanced to provide a more comprehensive view of your FSx for ONTAP file systems. The enhancements include new cards for SSD capacity, Capacity pool, and NetApp Autonomous Ransomware Protection with AI (ARP/AI). These cards summarize capacity utilization and ARP/AI protection for all volumes in the file system.

### **Support for second-generation Amazon FSx for NetApp ONTAP file systems**

Workload factory now supports second-generation Amazon FSx for NetApp ONTAP file systems. You can create, manage, and monitor second-generation file systems in the Workload Factory console. All new commercial regions are supported.

[Create a second-generation file system in Workload Factory](#)

### **FlexVol volume support for rebalancing volume capacity**

FlexVol volumes are discoverable within the Workload Factory console. You can check the balance of your FlexVol volumes and rebalance FlexVol volumes to redistribute the capacity when imbalances develop over time due to the addition of new files and file growth.

[Rebalance the capacity of a FlexVol volume](#)

## Terminology update

The term "Autonomous Ransomware Protection" (ARP) has been updated to "NetApp Autonomous Ransomware Protection with AI" (ARP/AI) in the Workload Factory console.

### ARP/AI enabled by default for new volumes

When you create a new volume in the Workload Factory console, NetApp Autonomous Ransomware Protection with AI (ARP/AI) is enabled by default if the file system has an ARP/AI policy. This means that the volume is automatically protected against ransomware attacks using AI-driven detection and response capabilities.

[Create a volume in Workload Factory](#)

### Replication support for immutable files

Workload factory supports replicating immutable volumes from one FSx for ONTAP system to another FSx for ONTAP file system to protect critical data from accidental deletion or malicious attacks like ransomware. The target volume and its host file system will be immutable, or locked, and any data in the target file system can't be modified or removed until the retention period ends.

[Learn how to create a replication relationship](#)

### Manage IAM execution role and permissions during link creation

Now you can manage the IAM execution role and its attached permission policy when you create a link in the Workload Factory console. A link establishes connectivity between your Workload Factory account and one or more FSx for ONTAP file systems. You have two options for assigning the IAM execution role and link permissions - automatically or user-provided. Managing the execution role and its attached permissions policy in Workload Factory means that you don't need to use third party code any longer.

[Connect to an FSx for ONTAP file system with a Lambda link](#)

## 08 June 2025

### New well-architected analysis and support for fixing issues

Automatic capacity management for FSx for ONTAP file systems is now included as a configuration analysis in the well-architected status dashboard.

Additionally, Workload Factory now supports fixing the following configuration issues:

- SSD capacity threshold
- Data tiering
- Scheduled local snapshots
- FSx for ONTAP backups
- Remote data replication
- Storage efficiencies
- Automatic capacity management

[Fix configuration issues](#)

## 03 June 2025

### Volume autogrow enhancement

Now you can set the autogrow size of your volumes so that volume size can grow beyond the provisioned size for business needs and application requirements.

[Enable volume autogrow](#)

### Well-architected analysis update

Workload factory now analyzes your FSx for ONTAP file systems to check whether storage efficiencies including data compaction, compression, and deduplication are being utilized. Storage efficiencies measure how effectively the file systems use available space.

[View the well-architected status of storage efficiencies](#)

### Storage dashboard enhancements

Starting today, when you open the Storage workload from the Workload Factory console, you'll view the **Dashboard**. The newly designed dashboard provides a holistic view of your FSx for ONTAP systems including the number of file systems, the total SSD capacity, the well-architected status overview, the data protection overview, and replication relationship health.

### Volumes tab enhancements

The Storage workload made enhancements to the Volumes tab within an FSx for ONTAP file system in the Workload Factory console. The enhancements include:

- **New cards:** SSD capacity, Capacity pool, and Autonomous Ransomware Protection (ARP)
- **New columns:** Capacity distribution, Used SSD capacity, Used capacity pool, and SSD efficiency

### Storage efficiencies update for volume creation

When creating a new volume, storage efficiencies including data compaction, compression, and deduplication are enabled by default.

[Create a new volume in Workload Factory](#)

## 04 May 2025

### Autonomous Ransomware Protection for FSx for ONTAP file systems

Protect your data with Autonomous Ransomware Protection (ARP), 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 also creates new, immutable snapshots from which you can restore your data.

[Protect your data with Autonomous Ransomware Protection](#)

### FlexGroup volume rebalance enhancement

BlueXP workload factory introduces the FlexGroup volume rebalance wizard with several layout options for rebalancing the data in a FlexGroup volume. Rebalancing redistributes data evenly to FlexGroup member

volumes.

[Rebalance the capacity in a FlexGroup volume](#)

### **Implement best practices for an FSx for ONTAP file system**

BlueXP workload factory provides a dashboard where you can review the well-architected status of your file system configurations. You can leverage this analysis to implement best practices for your FSx for ONTAP file systems. File system configuration analysis includes the following configurations: SSD capacity threshold, scheduled local snapshots, scheduled FSx for ONTAP backups, data tiering, and remote data replication.

- [Learn about the well-architected analysis for file system configurations](#)
- [Implement best practices for your file systems](#)

### **Dual-protocol volume security style options**

You have the option to choose either NTFS or UNIX as the security style for a volume to determine the method that users and permissions access a volume.

[Create a volume](#)

### **Replication enhancements**

#### **Reverse replication supported from FSx for ONTAP to on-premises**

Reverse replication is now available from an FSx for ONTAP file system to an on-premises ONTAP cluster from within the Workload Factory console.

[Reverse replication](#)

#### **Data protection volume replication**

You can now replicate data protection volumes.

[Replicate a data protection volume](#)

#### **Multiple volume selection**

Multiple volume selection is available so you can select exactly the volumes you want to replicate.

[Create a replication relationship](#)

#### **Long-term retention policy labels**

When you enable long-term retention for a replication relationship, source and target volumes labels must match exactly. Now BlueXP workload factory can automatically create missing source volume labels for you.

[Create a replication relationship](#)

### **FSx for ONTAP file name visible in volume creation**

We've improved the visibility of FSx for ONTAP file systems during volume creation. You'll see the FSx for ONTAP file system when you create a volume, so you'll know exactly where the volume is being created.

## **AWS account visible across the Storage workload**

We've improved account visibility across the Storage workload. You'll see the AWS account when navigating to the **Volumes**, **Storage VMs**, and **Replication** tabs.

## **Link association enhancements**

- You can quickly associate a link from an FSx for ONTAP file system in the Inventory tab.
- BlueXP workload factory now supports the use of alternative ONTAP user credentials for link association.

## **Link authentication support for AWS Secrets Manager**

You now have the option to use secrets from AWS Secrets Manager to authenticate links so that you don't have to use credentials stored in BlueXP workload factory.

## **Tracker response support**

Tracker now provides API responses so that you can see the REST API output related to the task.

[Monitor operations with Tracker](#)

## **Capacity validation when restoring a volume from a backup**

When restoring a volume from a backup, BlueXP workload factory determines if you have enough capacity for the restore and can automatically add SSD storage tier capacity if you don't.

[Restore a volume from a backup](#)

## **Support for alternative ONTAP user credentials**

Workload factory now supports alternative sets of ONTAP credentials for creating file systems to minimize security risks. Instead of using only the fsxadmin user, you can select a different set of ONTAP credentials or choose not to provide a password for fsxadmin and vsaadmin users.

## **Updated permissions terminology**

The Workload Factory user interface and documentation now use "read-only" to refer to read permissions and "read/write" to refer to automate permissions.

## **30 March 2025**

### **Automatic capacity management for scale-out systems**

Workload factory now scans for available inodes in volumes and increases their count according to the configured automatic capacity management thresholds. This feature supports automatic capacity management for scale-out systems. You can enable inodes management as part of automatic capacity management.

[Enable automatic capacity management](#)

### **FlexGroup rebalance API**

BlueXP workload factory releases the FlexGroup rebalance API that allows you to execute a plan to rebalance the data in a FlexGroup. Rebalancing redistributes data evenly to the member volumes.

## Replicate data form includes use cases

The replicate data form now includes use cases to make it easier for you to complete the form. You'll select one of the following use cases for data replication: migration, hot disaster recovery, cold disaster recovery, archive, or other. After you select a use case, Workload factory recommends values in accordance with best practices. You can accept the preselected values or customize the values in the form.

### [Replicate data](#)

## Data tiering policy terminology changes

Now when you select a tiering policy during volume creation, data replication, or updates to existing tiering policies, you'll find new terms to describe the tiering policies.

- *Balanced (Auto)*
- *Cost-optimized (All)*
- *Performance optimized (Snapshots only)*

## Security group details for file system creation

A security group is created as part of the FSx for ONTAP file system creation process. Security group details including protocols, ports, and roles are now available.

### [Create a file system](#)

## 02 March 2025

## Automatic capacity management improvements

When automatic capacity management is enabled, BlueXP workload factory now checks if a file system reached its capacity threshold every 30 minutes instead of every 2 hours.

The provisioned IOPS setting is no longer affected when capacity threshold is reached.

## Immutable snapshots

Now you can lock snapshots, making them immutable, for a specific retention period. Locking prevents the unauthorized access and malicious deletions of snapshots. You can enable immutable snapshots during snapshot policy creation, when creating manual snapshots, and after snapshot creation.

## Immutable files update

You can now make the following changes to your immutable files configuration: retention policy, retention period, autocommit period, and volume append mode.

### [Manage immutable files](#)

## Data replication enhancements

- **Cross-account replication:** Replication between two AWS accounts is supported in the BlueXP workload factory console as well as replication management.

- **Pause and resume replication:** You can pause (quiesce) scheduled replication updates from the source volume to the destination volume and then resume the replication schedule when you're ready. During the pause, source and destination volumes become independent, and the destination volume transitions from read-only to read/write.

[Pause and resume a replication relationship](#)

## CloudShell events in Tracker

Now you can track CloudShell events in Tracker.

[Learn how to monitor and track operations with Tracker](#)

## 02 February 2025

### CloudShell in BlueXP workload factory console

CloudShell is an embedded CLI capability available within BlueXP workload factory for Storage. You can use CloudShell to create, share, and execute ONTAP or AWS CLI commands from multiple sessions in a shell-like environment from within the Workload Factory console.

[Learn more about CloudShell in BlueXP workload factory](#)

### Inventory data download

You can now download FSx for ONTAP inventory data into an Microsoft Excel or CSV file from Storage in BlueXP workload factory.

Name	Status	AWS account	Region	SSD storage size	Capacity pool size	Tags	Creation time
fsx-wimdb-DEFAULT	AVAILABLE	627023167428	US East (N. Virginia)   us-east-1	2 TiB	574.66 GiB	1 View	Jan 27, 2025, 9:13 PM

### FSx for ONTAP file system additional menu options

We've made it simpler to do the following for an FSx for ONTAP file system from the FSx for ONTAP tab in Storage.

- Create a storage VM
- Create a volume
- Replicate volume data

### Terraform support for creating volumes

You can now use Terraform from the Codebox to create volumes.

[Create a volume](#)

### File locking with the immutable files feature

You can now lock files using the immutable files feature when you create a volume for an FSx for ONTAP file

system. File locking helps you and others prevent accidental or intentional file deletion for a specified period.

[Create a volume](#)

### **Tracker available for monitoring and tracking operations**

Tracker, a new monitoring capability is available in Storage. You can use Tracker to monitor and track the progress and status of credentials, storage, and link operations, review details for operation tasks and subtasks, diagnose any issues or failures, edit parameters for failed operations, and retry failed operations.

[Learn how to monitor and track operations with Tracker](#)

### **Support for second-generation Amazon FSx for NetApp ONTAP file systems**

You can now use Amazon FSx for NetApp ONTAP second-generation file systems in NetApp Workload Factory. FSx for ONTAP second-generation Single-AZ file systems are powered by up to 12 HA pairs which can deliver up to 72 GBps of throughput capacity and 2,400,000 SSD IOPS. FSx for ONTAP second-generation Multi-AZ file systems are powered by one HA pair and deliver 6 GBps of throughput capacity and 200,000 SSD IOPS.

- [Add high-availability pairs](#)
- [Quotas and limits for Amazon FSx for NetApp ONTAP](#)

## **05 January 2025**

### **Volume CIFS share enhancements**

The following enhancements are available for managing CIFS share for volumes in an Amazon FSx for ONTAP file system in BlueXP workload factory:

- Support for multiple CIFS shares on a volume
- The option to update users and groups at any time
- The option to update permissions for users and groups at any time
- CIFS share deletion

[Manage CIFS shares](#)

## **1 December 2024**

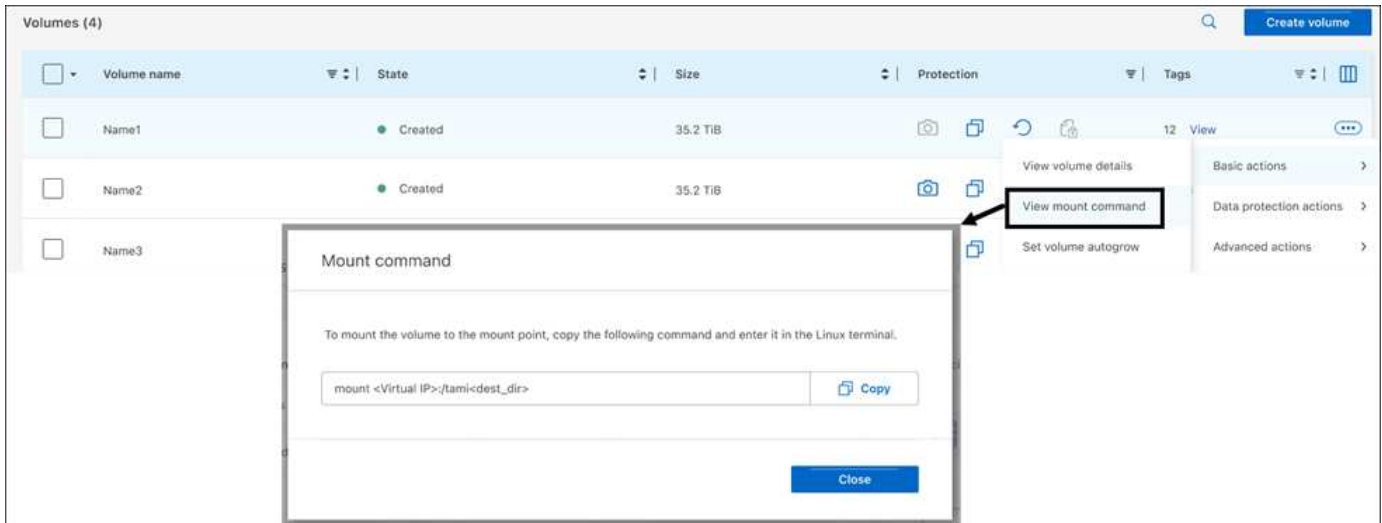
### **Block storage for scale-out FSx for ONTAP file systems**

You can now provision block storage over FSx for ONTAP when using a scale-out file system deployment with up to 6 HA pairs.

[Create an FSx for ONTAP file system in BlueXP workload factory](#)

### **Mount command available**

Mount commands are now available for NFS and CIFS access to a volume. You can get the mount point for a volume within an FSx for ONTAP file system by selecting **Basic actions** then **View mount command**.



[View mount command for a volume](#)

### Update storage efficiency after volume creation

You can now enable or disable storage efficiency for FlexVol volumes after volume creation. Storage efficiency includes deduplication, data compression, and data compaction. Enabling storage efficiency helps you achieve optimal space savings on a FlexVol volume.

[Update storage efficiency for a volume](#)

### On-premises ONTAP cluster discovery and replication

Discover and replicate on-premises ONTAP cluster data to an FSx for ONTAP file system so that it can be used to enrich AI knowledge bases. All on-premises discovery and replication workflows are possible from the new **On-Premises ONTAP** tab in the Storage inventory.

[Discover an on-premises ONTAP cluster](#)

### AWS credentials improve savings calculator analysis

You now have the option to add AWS credentials from the savings calculator. Adding credentials improves the accuracy of the savings calculator analysis of your Amazon Elastic Block Store, Elastic File Systems, and FSx for Windows File Server storage environments when compared with FSx for ONTAP.

[Explore savings with FSx for ONTAP in BlueXP workload factory](#)

## 3 November 2024

### Tab views in storage inventory

Storage inventory has been updated to a two-tab view:

- FSx for ONTAP tab: displays the FSx for ONTAP file systems you currently have.
- Explore savings tab: displays Elastic Block Store, FSx for Windows File Server, and Elastic File Systems storage systems. From there, you can explore savings for these systems by comparing them with FSx for ONTAP.

## 29 September 2024

### Link creation updates

- **Codebox viewer:** Codebox is now integrated in the link creation process. You can view and copy the CloudFormation template from Codebox in Workload Factory before redirecting to AWS to execute the operation.
- **Required permissions:** The permissions required to execute the link creation in AWS CloudFormation are now available to view and copy from the Create Link wizard in Workload Factory.
- **Support for manual link creation:** This feature allows standalone creation in AWS CloudFormation with manual registration of the link ARN. It's useful when a Security or DevOps team assists in the link creation process.

[Create a link](#)

## 1 September 2024

### Read mode support for storage management

Read mode is available for storage management in Workload Factory. Read mode enhances the experience of basic mode by adding read-only permissions so that the Infrastructure-as-Code templates are filled with your specific variables. The Infrastructure-as-Code templates can be executed directly from your AWS account without providing any modify permissions to Workload Factory.

### Backup before volume deletion support

You can now back up a volume before deleting it. The backup will remain in the file system until deleted.

[Delete a volume](#)

## 4 August 2024

### Terraform support

You can now use Terraform from the Codebox to deploy file systems and storage VMs.

- [Create a file system](#)
- [Create a storage VM](#)
- [Use Terraform from Codebox](#)

### Throughput and IOPS recommendations in the storage calculator

The storage calculator makes FSx for ONTAP file system configuration recommendations for throughput and IOPS based on AWS best practices, which provides you with optimal guidance for your selections.

## 7 July 2024

### Initial release of Workload Factory for Amazon FSx for NetApp ONTAP

Amazon FSx for NetApp ONTAP is now generally available in BlueXP workload factory.

# Known limitations of Amazon FSx for NetApp ONTAP in NetApp Workload Factory

Known limitations identify platforms, devices, or functions that are not supported by this release of the product, or that do not interoperate correctly with it. The following limitations are specific to Amazon FSx for NetApp ONTAP in NetApp Workload Factory. Review these limitations carefully.

## Region support

The following AWS regions are not supported:

- China regions
- GovCloud (US) regions
- Secret Cloud
- Top Secret Cloud

## Adding HA pairs limitations

- This operation might take approximately 30 minutes to complete.
- Adding HA pairs limits the following operations: adding more HA pairs, provisioning IOPS, increasing SSD capacity, and updating throughput capacity.

## Throughput capacity region support

### Scale-up deployments

Scale-up configurations are supported up to 2 GB/s in all regions. The following regions support up to 4 GB/s maximum throughput capacity for scale-up deployments: US East (Ohio) Region, US East (N. Virginia) Region, US West (Oregon) Region, and Europe (Ireland).

### Scale-out deployments

The following regions support up to 6 GB/s maximum throughput capacity for scale-out deployments: US East (N. Virginia and Ohio), US West (N. California and Oregon), Europe (Frankfurt, Ireland, and Stockholm), and Asia Pacific (Singapore and Sydney).

## Capacity management

- The volume autogrow feature isn't currently supported for the iSCSI protocol.
- When the automatic capacity management feature is running, manual capacity increase isn't allowed.
- Disabling the automatic capacity management feature is only possible with the same permissions used to enable it.
- When automatic capacity management is enabled, a link is required to make sure volume inodes increase along with storage capacity.

## Storage VMs

The number of storage VMs is limited per SKU. Creating storage VMs beyond the limitation isn't supported in Workload Factory.<sup>1</sup>

Refer to [Managing FSx for ONTAP storage virtual machines](#) in AWS documentation for the maximum number of SVMs per file system.

## iSCSI protocol support

- The iSCSI protocol is only available for FlexVol volumes.<sup>1</sup>
- Volume size decreases aren't supported for iSCSI volumes.

## Data protection

- Snapshots cannot be deleted.
- When you replicate a file system, all volumes in the file system use the same replication policy.
- For long-term retention replication relationships, only the last snapshot is available for restore.
- The following features aren't supported with immutable snapshots:
  - Consistency groups
  - FabricPool
  - FlexCache volumes
  - SMtape
  - SnapMirror active sync
  - SnapMirror policy rules using the `-schedule` parameter
  - SnapMirror synchronous
  - SVM data mobility (used for migrating or relocating an SVM from a source cluster to a destination cluster)

## Storage savings calculator

The Storage savings calculator doesn't calculate cost savings for the following configurations:

- FSx for Windows File Server: HDD storage type
- Elastic Block Store (EBS): `st1`, `sc1`, and standard volume types
- Elastic File System (EFS): Bursting throughput mode

## AWS Secrets Manager support

AWS Secrets Manager isn't supported when using a Console agent.

## Amazon S3 access points limitation

The limit for the number of S3 access points per storage VM is 4,000.

Note:

1. Applies to Amazon FSx for NetApp ONTAP

# Get started

## Learn about Amazon FSx for NetApp ONTAP in NetApp Workload Factory

Amazon FSx for NetApp ONTAP is a fully managed, cloud-based data storage service that provides advanced data management capabilities and highly scalable performance. FSx for ONTAP allows you to create and manage file systems as the storage backend for all your workloads within NetApp Workload Factory.

FSx for ONTAP provides the same features, performance, and administrative capabilities NetApp customers use on premises today, with the simplicity, agility, security, and scalability of a native AWS service.

FSx for ONTAP is the *Storage* component in Workload Factory.

### Features

FSx for ONTAP offers the following features:

- **Fully-managed service:** provides a fully-managed service integrated with the Workload Factory console.
- **High availability:** provides high availability for each FSx for ONTAP file system, supporting Single and Multiple Availability Zones deployments.
- **Automated snapshots:** protects data with automated, efficient snapshots, which are near instantaneous, space efficient point-in-time read-only copies of the file system or volumes.
- **Volume replication:** provides disaster recovery with cross-region replication across Amazon Web Services.
- **Efficient backups:** adds an extra layer of protection with a copy of the data in another region for emergencies.
- **Fast cloning:** accelerates application development with fast cloning.
- **Multi-protocol support:** supports Network File System (NFS), Server Message Block (SMB), and Internet Small Computer Systems Interface (iSCSI) protocols.
- **High throughput:** delivers high throughput performance to ensure low latencies for workloads running on top of FSx for ONTAP file systems.
- **In-memory cache and NVMe cache:** includes a unique in-memory cache and NVMe cache, which boosts the performance of frequently accessed data.
- **Hundreds of thousands of IOPS:** provides hundreds of thousands of IOPS with SSD disks, ensuring that your storage and workloads receive timely results.
- **Thin Provisioning:** allows capacity provisioning in advance, saving costs until more capacity is needed.
- **Data deduplication and compression:** removes duplicate data and compresses data to reduce the amount of physical storage that is required for FSx for ONTAP file systems resulting in cost savings.
- **Data tiering:** allows storage cost reduction by moving less frequently accessed data from the primary, high performance SSD storage tier to the secondary capacity pool storage tier.

## Additional features in Workload Factory

- **Storage cost comparison calculator:** compares your Amazon Elastic Block Store (EBS), Elastic File System (EFS) and FSx for Windows File Server storage costs with FSx for ONTAP. From the calculator, you can view how FSx for ONTAP storage configurations offer potential savings and plan your move to FSx for ONTAP storage.
- **Workload Factory user interface:** provides *Quick create* and *Advanced create* deployment mode options. Quick create includes AWS, NetApp, and industry standard best practices for your storage configurations.
- **Codebox:** provides developers with a code viewer for FSx for ONTAP operations, code templates for copy and download, and an automation catalog for code re-use.

## Tools to use NetApp Workload Factory

You can use NetApp Workload Factory with the following tools:

- **Workload Factory console:** The Workload Factory console provides a visual, holistic view of your applications and projects.
- **NetApp Console:** The NetApp Console provides a hybrid interface experience so that you can use Workload Factory along with other NetApp data services.
- **Ask me:** Use the Ask me AI assistant to ask questions and learn more about Workload Factory without leaving the Workload Factory console. Access Ask me from the Workload Factory help menu.
- **CloudShell CLI:** Workload Factory includes a CloudShell CLI to manage and operate AWS and NetApp environments across accounts from a single, browser-based CLI. Access CloudShell from the top bar of the Workload Factory console.
- **REST API:** Use the Workload Factory REST APIs to deploy and manage your FSx for ONTAP file systems and other AWS resources.
- **CloudFormation:** Use AWS CloudFormation code to perform the actions you defined in the Workload Factory console to model, provision, and manage AWS and third-party resources from the CloudFormation stack in your AWS account.
- **Terraform NetApp Workload Factory provider:** Use Terraform to build and manage infrastructure workflows generated in the Workload Factory console.

## Cost

AWS maintains your FSx for ONTAP account, not Workload Factory. Refer to [Pricing for Amazon FSx for NetApp ONTAP](#).

## Regions

Workload factory is supported in all commercial regions where FSx for ONTAP is supported. [View supported Amazon regions](#).

The following AWS regions aren't supported:

- China regions
- GovCloud (US) regions
- Secret Cloud
- Top Secret Cloud

## Getting help

Amazon FSx for NetApp ONTAP is an AWS first-party solution. For questions or technical support issues associated with your FSx for ONTAP file system, infrastructure, or any solution using this service, use the Support Center in your AWS Management Console to open a support case with AWS. Select the “FSx for ONTAP” service and appropriate category. Provide the remaining information required to create your AWS support case.

For general questions about Workload Factory or Workload Factory applications and services, refer to [Get help for FSx for ONTAP for Workload Factory](#).

## Quick start for Amazon FSx for NetApp ONTAP in NetApp Workload Factory

With Amazon FSx for NetApp ONTAP in NetApp Workload Factory, you can get started immediately in *basic* mode.

If you'd like to use Workload Factory to create a file system, manage resources, and more, you can get started in a few steps. In this case, you need an AWS account and credentials to get started.

Follow these steps to get started.

1

### Log in to Workload Factory

You'll need to [set up an account with Workload Factory](#) and [log in](#)

2

### Add credentials and permissions

Choose the [permission policies](#) to meet your needs.

If you choose not to grant permissions, you can start using Workload Factory for FSx for ONTAP to copy partially completed code samples.

If you choose to grant permissions, you'll need to [add credentials to an account manually](#) that includes selecting workload capabilities, such as Databases and AI, and creating the IAM policies for the required permissions.

3

### Create a file system

You'll create an FSx file system to begin managing your storage and FSx for ONTAP resources in Workload Factory. In the [Workload Factory console](#), in Storage, select **Create file system**. [Learn how to create a file system](#).

You can also start with the storage savings calculator to compare the costs of your Amazon Elastic Block Store, Elastic File System, and FSx for Windows File Server storage environments to that of FSx for ONTAP. [Explore savings with the storage savings calculator](#).

### What's next

With a file system in your Storage inventory, you can [create volumes](#), manage your FSx for ONTAP file system, and set up data protection for your resources.

# Create an FSx for ONTAP file system in NetApp Workload Factory

Using NetApp Workload Factory you can create first and second-generation FSx for ONTAP file systems to add and manage volumes and additional data services.

## About this task

A storage VM and a security group are created as part of file system creation.

## Before you begin

Before creating your FSx for ONTAP file system, you will need:

- Credentials with *file system creation and deletion* permissions to create an FSx for ONTAP file system. [Learn how to grant permissions to an AWS account.](#)
- The region and VPC information for where you will create the FSx for ONTAP instance.

## Create an FSx for ONTAP file system

You can create an FSx for ONTAP file system using *Quick create* or *Advanced create*. You can also use the following tools available in the Codebox: REST API, CloudFormation, and Terraform. [Learn how to use Codebox for automation.](#)



When using Terraform from Codebox, the code you copy or download hides `fsxadmin` and `vsadmin` passwords. You'll need to re-enter the passwords when you run the code.

## Quick create

Quick create enables you to use a recommended best-practice configuration. You can change most settings after you create an FSx for ONTAP file system.

Second-generation FSx for ONTAP file systems are the default deployment type for quick create unless the selected region doesn't support second-generation FSx for ONTAP file systems.

## Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage dashboard, select **Create file system**.
4. On the Create FSx for ONTAP file system page, select **Quick create**.

You can also load a saved configuration.

5. Under File system general configuration, provide the following:
  - a. **AWS credentials:** Select to add AWS credentials in Workload Factory or continue without credentials.
  - b. **File system name:** Enter a name for the file system.
  - c. **Region & VPC:** Select the region and VPC for the file system.
  - d. **Deployment type:** Select a deployment type.

- **Single Availability Zone (Single-AZ) deployment:** provides availability by monitoring for hardware failures and automatically replacing infrastructure components in the event of a failure. Achieves high durability by automatically replicating your data within an Availability Zone to protect it from component failure.

This configuration is recommended for high performance workloads or when workloads start small and incrementally scale out to 72 GB/s of throughput and 2.4 million IOPS.

- **Multiple Availability Zones (Multi-AZ) deployment:** provides continuous availability to data even when an Availability Zone is unavailable. A Multi-AZ file system is designed for business-critical production workloads that require high availability to shared ONTAP file data and need storage with built-in replication across Availability Zones.

This single HA-pair configuration is recommended for workloads that require up to 6 GB/s of throughput or 200,000 IOPS.

- e. **Tags:** Optionally, you can add up to 50 tags.
6. Under **File system details**, provide the following:
    - a. **SSD storage capacity:** Enter the storage capacity and select the storage capacity unit.
      - For first-generation deployments, you can't decrease capacity after file system creation.
      - For second-generation deployments, you can increase capacity after file system creation.
    - b. **ONTAP credentials:** Optional. Enter your ONTAP user name and password. The password can be set now or later.

If the user you provide is not the fsxadmin user, and later you need to reset the fsxadmin password, you'll be able to do this from the AWS console.

- c. **SMB/CIFS setup:** Optional. If you plan to use SMB/CIFS protocol to access volumes, you must configure the Active Directory for the storage VM during file system creation. Provide the following details for the storage VM that is created for this file system.
- i. **Active Directory domain to join:** Enter the fully qualified domain name (FQDN) for the Active Directory.
  - ii. **DNS IP addresses:** Enter up to three DNS IP addresses separated by commas.
  - iii. **SMB server NetBIOS name:** Enter the SMB server NetBIOS name of the Active Directory computer object to create for your storage VM. This is the name of this storage VM in the Active Directory.
  - iv. **User name:** Enter the user name of the service account in your existing Active Directory.

Do not include a domain prefix or suffix. For `EXAMPLE\ADMIN`, use `ADMIN`.

- v. **Password:** Enter the password for the service account.
- vi. **Organization unit:** Optionally, enter the name of the Organizational Unit where you intend to create the computer account for FSx for ONTAP. The OU is the distinguished path name of the organizational unit to which you want to join the file system.
- vii. **Delegated administrators group:** Optionally, enter the name of the group in your Active Directory that can administer your file system.

If you are using AWS Managed Microsoft AD, you must specify a group such as `AWS Delegated FSx Administrators`, `AWS Delegated Administrators`, or a custom group with delegated permissions to the OU.

If you are joining to a self-managed AD, use the name of the group in your AD. The default group is `Domain Admins`.

7. Open the **Summary** to review the configuration that you defined. If needed, you can change any setting at this time before saving or creating the file system.
8. Save or create the file system.

If you created the file system, you can now view the FSx for ONTAP file system in the **Inventory** page.

### Advanced create

With Advanced create, you set all of the configuration options, including availability, security, backups, and maintenance.

### Steps

1. Log in using one of the [console experiences](#).
2. In the Storage tile, select **Create FSx for ONTAP**.
3. On the Create FSx for ONTAP file system page, select **Advanced create**.

You can also load a saved configuration.

4. Under File system general configuration, provide the following:
  - a. **AWS credentials:** Select to add AWS credentials in Workload Factory or continue without credentials.
  - b. **File system name:** Enter a name for the file system.

- c. **Region & VPC:** Select the region and VPC for the file system.
- d. **Deployment type:** Select a deployment type and file system generation. The availability of a second-generation file system depends on the selected region. If the selected region doesn't support second-generation FSx for ONTAP file systems, the deployment type switches to first-generation.

- **Single Availability Zone (Single-AZ) deployment:** provides availability by monitoring for hardware failures and automatically replacing infrastructure components in the event of a failure. Achieves high durability by automatically replicating your data within an Availability Zone to protect it from component failure.

**File system generation:** Select one of the following:

- **Second-generation:** This configuration is recommended for high performance workloads or when workloads start small and incrementally scale out to 72 GB/s of throughput and 2.4 million IOPS.
- **First-generation:** This configuration is ideal for workloads that require up to 4 GB/s or 160,000 IOPS. First-generation file systems can only increase capacity.
- **Multiple Availability Zones (Multi-AZ) deployment:** provides continuous availability to data even when an Availability Zone is unavailable. A Multi-AZ file system is designed for business-critical production workloads that require high availability to shared ONTAP file data and need storage with built-in replication across Availability Zones.

**File system generation:** Select one of the following:

- **Second-generation:** This single HA-pair configuration is recommended for workloads that require up to 6 GB/s of throughput or 200,000 IOPS. In a Multi-AZ and second-generation file system, capacity can increase or decrease to match workload demands.
- **First-generation:** This configuration is ideal for workloads that require up to 4 GB/s or 160,000 IOPS. First-generation file systems can only increase capacity.

- e. **Tags:** Optionally, you can add up to 50 tags.

5. Under File system details, provide the following:

- a. **SSD storage capacity:** Enter the storage capacity and select the storage capacity unit.
- For first-generation deployments, you can't decrease capacity after file system creation.
  - For second-generation deployments, you can adjust capacity.
- b. **Throughput capacity per HA pair:** Select throughput capacity per number of HA pairs. First-generation file systems support only one HA pair.
- c. **Provisioned IOPS:** Select one of the following options:
- **Automatic:** For automatic, for every GiB created, 3 IOPS are added.
  - **User-provisioned:** For user-provisioned, enter the IOPS value.
- d. **ONTAP credentials:** Optional. Enter your ONTAP user name and password. The password can be set now or later.

If the user you provide is not the fsxadmin user, and later you need to reset the fsxadmin password, you'll be able to do this from the AWS console.

- e. **Storage VM Credentials:** Optional. Enter your user name. Password can be specific to this file system or you can use the same password entered for ONTAP credentials. The password can be set now or later.

- f. **SMB/CIFS setup:** Optional. If you plan to use SMB/CIFS protocol to access volumes, you must configure the Active Directory for the storage VM during file system creation. Provide the following details for the storage VM that is created for this file system.
- i. **Active Directory domain to join:** Enter the fully qualified domain name (FQDN) for the Active Directory.
  - ii. **DNS IP addresses:** Enter up to three DNS IP addresses separated by commas.
  - iii. **SMB server NetBIOS name:** Enter the SMB server NetBIOS name of the Active Directory computer object to create for your storage VM. This is the name of this storage VM in the Active Directory.
  - iv. **User name:** Enter the user name of the service account in your existing Active Directory.

Do not include a domain prefix or suffix. For `EXAMPLE\ADMIN`, use `ADMIN`.

- v. **Password:** Enter the password for the service account.
- vi. **Organization unit:** Optionally, enter the name of the Organizational Unit where you intend to create the computer account for FSx for ONTAP. The OU is the distinguished path name of the organizational unit to which you want to join the file system.
- vii. **Delegated administrators group:** Optionally, enter the name of the group in your Active Directory that can administer your file system.

If you are using AWS Managed Microsoft AD, you must specify a group such as `AWS Delegated FSx Administrators`, `AWS Delegated Administrators`, or a custom group with delegated permissions to the OU.

If you are joining to a self-managed AD, use the name of the group in your AD. The default group is `Domain Admins`.

6. Under Network & security, provide the following:

- a. **Security group:** Create or use an existing security group.

For a new security group, refer to [security group details](#) for a description of the security group protocols, ports, and roles.

- b. **Availability Zones:** Select availability zones and subnets.
  - For Cluster configuration node 1: Select an availability zone and subnet.
  - For Cluster configuration node 2: Select an availability zone and subnet.
- c. **VPC route tables:** Select the VPC route table to enable client access to volumes.
- d. **Endpoint IP address range:** Select **Floating IP address range outside your VPC** or **Enter an IP address range** and enter an IP address range.
- e. **Encryption:** Select the encryption key name from the dropdown.

7. Under Backup and maintenance, provide the following:

- a. **Volume backups:** Daily automatic backups are enabled by default. Disable if desired.
  - i. **Automatic backup retention period:** Enter the number of days to retain automatic backups.
  - ii. **Daily automatic backup window:** Select either **No preference** (a daily backup start time is selected for you) or **Select start time for daily backups** and specify a start time.
- b. **Weekly maintenance window:** Select either **No preference** (a weekly maintenance window start time is selected for you) or **Select start time for 30-minute weekly maintenance window** and

specify a start time.

8. Save or create the file system.

If you created the file system, you can now view the FSx for ONTAP file system in the **Inventory** page.

## Security group details

The following table provides security group details including protocols, ports, and roles.

<b>Protocol</b>	<b>Port</b>	<b>Role</b>
SSH	22	SSH access to the IP address of the cluster management LIF or a node management LIF
TCP	80	Web page access to the IP address of the cluster management LIF
TCP/UDP	111	Remote procedure call for NFS
TCP/UDP	135	Remote procedure call for CIFS
UDP	137	NetBIOS name resolution for CIFS
TCP/UDP	139	NetBIOS service session for CIFS
TCP	443	ONTAP REST API access to the IP address of the cluster management LIF or an SVM management LIF
TCP	445	Microsoft SMB/CIFS over TCP with NetBIOS framing
TCP/UDP	635	NFS mount
TCP	749	Kerberos
TCP/UDP	2049	NFS server daemon
TCP	3260	iSCSI access through the iSCSI data LIF
TCP/UDP	4045	NFS lock daemon
TCP/UDP	4046	Network status monitor for NFS
UDP	4049	NFS quota protocol
TCP	10000	Network data management protocol (NDMP) and NetApp SnapMirror intercluster communication
TCP	11104	Management of NetApp SnapMirror intercluster communication
TCP	11105	SnapMirror data transfer using intercluster LIFs

Protocol	Port	Role
TCP/UDP	161-162	Simple network management protocol (SNMP)
All ICMP	All	Pinging the instance

**What's next**

With a file system in your Storage inventory, you can [create volumes](#), manage your FSx for ONTAP file system, and set up [data protection](#) for your resources.

# Use Amazon FSx for NetApp ONTAP

## Explore savings with FSx for ONTAP in NetApp Workload Factory

Explore savings for your storage workloads that use Amazon Elastic Block Store (EBS), Elastic File System (EFS), and FSx for Windows File Server against FSx for NetApp ONTAP.

NetApp Workload Factory includes a storage savings calculator to compare Amazon storage environments to FSx for ONTAP. You can explore savings with or without providing your AWS credentials and customize configuration settings for your storage environment. When you provide AWS credentials, you can select one or more instances of Amazon Elastic Block Store, for example, and let Workload Factory make the comparison automatically. Whether manually or automatically, the calculator determines which storage service provides the lowest cost for your storage needs.

If the storage calculator determines that the most cost-effective storage is FSx for ONTAP, you can create or save FSx for ONTAP configurations and use the Codebox to generate Infrastructure-as-Code templates regardless of the permissions you grant to Workload Factory.

### Calculator options

Two calculator options are available for making the cost comparison between your systems and FSx for ONTAP — customization and automatic detection for your Amazon storage environments.

Explore savings via customization: You provide the configuration settings for a storage environment including the use case, region, number of volumes or file systems, storage amount, snapshot frequency, amount changed per snapshot, provisioned IOPS, throughput, and more.

Explore savings for detected storage environments: Workload Factory links to your existing AWS storage environments and pulls in the details to the calculator for automatic comparison. You'll need to grant automate permissions to use automatic mode. You can change the use case, but all other details are automatically determined in the calculation.

Additionally, you can [add AWS credentials](#) to improve the accuracy of the calculator analysis. Select **Calculate savings based on existing resources**. You'll be redirected to the Add credentials page. After you add credentials, select the existing resources to compare with FSx for ONTAP, and select **Explore savings**.

### Explore savings via customization

Follow the steps under the tab for your storage selection.

## Amazon Elastic Block Store (EBS)

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **Explore savings**, and then select the **Amazon Elastic Block Store (EBS)** tab.
4. In the Storage savings calculator, provide the following details:
  - a. **Use case:** Required. Select a use case from the dropdown menu. The selected use case determines the FSx for ONTAP file system characteristics for comparison.
  - b. **Region:** Optional. Select the region for your EBS configuration from the dropdown menu.
  - c. **Select EBS volume type:** Optional. Select the EBS volume type used for your configuration.
  - d. **Number of volumes:** Optional. Enter the number of volumes in your EBS configuration.
  - e. **Storage amount per volume (TiB):** Optional. Enter the storage amount per volume in TiB.
  - f. **Snapshot frequency:** Optional. Select the snapshot frequency for your EBS configuration.
  - g. **Amount changed per snapshot (GiB):** Optional. For snapshot storage only. Enter the amount changed per snapshot in GiB.
  - h. **Provisioned IOPS per volume:** Optional. For gp3, io1, and io2 volumes. Enter the provisioned IOPS per volume.
  - i. **Throughput (MiB/s):** Optional. For gp3 volumes only. Enter throughput in MiB/s per volume.

## Amazon FSx for Windows File Server

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **Explore savings**, and then select the **Amazon Elastic Block Store (EBS)** tab.
4. In the Storage savings calculator, provide the following details:
  - a. **Use case:** Required. Select a use case from the dropdown menu. The selected use case determines the FSx for ONTAP file system characteristics for comparison.
  - b. **Region:** Optional. Select the region for your FSx for Windows File Server configuration from the dropdown menu.
  - c. **Deployment type:** Optional. Select **Single Availability Zone** or **Multiple Availability Zones**.
  - d. **Storage type:** SSD storage type is selected by default.
  - e. **Storage capacity (TiB):** Optional. Enter the storage capacity for the configuration.
  - f. **Deduplication savings (%):** Optional. Enter the capacity savings percentage you expect from deduplication.
  - g. **Snapshot frequency:** Optional. Select the snapshot frequency for your configuration.
  - h. **Amount changed per snapshot (GiB):** Optional. For snapshot storage only. Enter the amount changed per snapshot in GiB.
  - i. **Provisioned SSD IOPS:** Optional. Enter the provisioned SSD IOPS.
  - j. **Throughput (MiB/s):** Optional. Enter throughput in MiB/s.

## Amazon Elastic File System (EFS)

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **Explore savings**, and then select the **Amazon Elastic Block Store (EBS)** tab.
4. In the Storage savings calculator, provide the following details:
  - a. **Use case:** Required. Select a use case from the dropdown menu. The selected use case determines the FSx for ONTAP file system characteristics for comparison.
  - b. **Region:** Optional. Select the region for your FSx for Windows File Server configuration from the dropdown menu.
  - c. **File System Type:** Optional. Select **Regional** or **One zone**.
  - d. **Storage capacity (TiB):** Optional. Enter the storage capacity of the EFS configuration.
  - e. **Data frequently accessed (%):** Optional. Enter the percentage of data that is frequently accessed.
  - f. **Throughput mode:** Optional. Select **Provisioned throughput** or **Elastic throughput**.
  - g. **Throughput (MiB/s):** Optional. Enter the throughput in MiB/s.

After you provide details for your storage system configuration, review the calculations and recommendations provided on the page.

Additionally, scroll down to the bottom of the page to view the report by selecting one of the following:

- **Export PDF**
- **Send by email**
- **View the calculations**

To switch to FSx for ONTAP, follow the instructions to [deploy FSx for ONTAP file systems](#).

## Explore savings for detected storage environments

### Before you begin

For Workload Factory to detect Amazon Elastic Block Store (EBS), Elastic File System (EFS), and FSx for Windows File Server storage environments in your AWS account, make sure you [grant view, planning, and analysis permissions](#) in your AWS account.



This calculator option doesn't support calculations for EBS snapshots and FSx for Windows File Server shadow copies. When exploring savings via customization, you can provide EBS and FSx for Windows File Server snapshot details.

Follow the steps under the tab for your storage selection.

## Amazon Elastic Block Store (EBS)

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **Explore savings**, and then select the **Amazon Elastic Block Store (EBS)** tab.
4. In the **Elastic Block Store (EBS)** tab, select the instance(s) to compare with FSx for ONTAP and select **Explore savings**.
5. The Storage savings calculator appears. The following storage system characteristics are pre-filled based on the instance(s) you selected:
  - a. **Use case:** The use case for your configuration. You can change the use case if needed.
  - b. **Selected volumes:** the number of volumes in the EBS configuration
  - c. **Total storage amount (TiB):** the storage amount per volume in TiB
  - d. **Total provisioned IOPS:** for gp3, io1, and io2 volumes
  - e. **Total throughput (MiB/s):** for gp3 volumes only

## Amazon FSx for Windows File Server

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **Explore savings**, and then select the **Amazon FSx for Windows File Server** tab.
4. In the **Amazon FSx for Windows File Server** tab, select the instance(s) to compare with FSx for ONTAP and select **Explore savings**.
5. The Storage savings calculator appears. The following storage system characteristics are pre-filled based on the deployment type of the instance(s) you selected:
  - a. **Use case:** The use case for your configuration. You can change the use case if needed.
  - b. **Selected file systems**
  - c. **Total storage amount (TiB)**
  - d. **Provisioned SSD IOPS**
  - e. **Throughput (MiB/s)**

## Amazon Elastic File System (EFS)

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **Explore savings**, and then select the **Amazon Elastic File System (EFS)** tab.
4. In the **Elastic File System (EFS)** tab, select the instance(s) to compare with FSx for ONTAP and select **Explore savings**.
5. The Storage savings calculator appears. The following storage system characteristics are pre-filled based on the instance(s) you selected:

- a. **Use case:** The use case for your configuration. You can change the use case if needed.
- b. **Total file systems**
- c. **Total storage amount (TiB)**
- d. **Total provisioned throughput (MiB/s)**
- e. **Total elastic throughput - read (GiB)**
- f. **Total elastic throughput – write (GiB)**

After you provide details for your storage system configuration, review the calculations and recommendations provided on the page.

Additionally, scroll down to the bottom of the page to view the report by selecting one of the following:

- **Export PDF**
- **Send by email**
- **View the calculations**

## Deploy FSx for ONTAP file systems

If you'd like to switch to FSx for ONTAP to realize cost savings, select **Create** to create the file system(s) directly from the Create an FSx for ONTAP file system wizard or select **Save** to save the recommended configuration(s) for later.

### Deployment methods

In *automate* mode, you can deploy the FSx for ONTAP file system directly from Workload Factory. You can also copy the content from the Codebox window and deploy the system using one of the Codebox methods.

In *basic* mode, you can copy the content from the Codebox window and deploy the FSx for ONTAP file system using one of the Codebox methods.

## Track costs for your resources in NetApp Workload Factory

Use NetApp Workload Factory to track FSx for ONTAP file system costs and usage in a consolidated view. The cost data helps you manage budgets and optimize resources effectively. AWS Cost Explorer provides the cost data.

### About this task

Cost and usage data for your FSx for ONTAP file system resources is extracted from AWS Cost Explorer using the following permissions:

- `ce:GetCostAndUsage`
- `ce:GetTags`

### Before you begin

[Grant credentials with the \*view, planning, and analysis\* permission policy](#) in Workload Factory to track FSx for ONTAP costs.

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **Cost**.
4. On the Cost page, filter cost and capacity data for your FSx for ONTAP resources by providing the following:
  - a. **AWS accounts**: Select the accounts you want to view cost data for.
  - b. **Credentials**: Select credentials with *view, planning, and analysis* permissions.
  - c. **Regions**: Select the AWS regions you want to view cost data for.
  - d. **Cost data range**: Select the time range for cost data to view.
5. View the **Cost Details** for your FSx for ONTAP resources.

## Use links

### Learn about NetApp Workload Factory links

A NetApp Workload Factory link creates a trust relationship and connectivity between a Workload Factory account and one or more FSx for ONTAP file systems. This allows you to monitor and manage certain file system features directly from the ONTAP REST API calls that are not available through the Amazon FSx for ONTAP API.

You don't need a link to get started with Workload Factory, but in some cases you'll need to create a link to unlock all Workload Factory features and workload capabilities.

#### Why links are beneficial

Links are beneficial because they allow Workload Factory to perform operations that are not natively available through the Amazon FSx for ONTAP API. Links enable advanced ONTAP capabilities and automations, which enhance the management of FSx for ONTAP file systems.

Here are some benefits of using links:

- The link enables the NetApp Console to send ONTAP commands directly to your FSx for ONTAP file system, bringing advanced ONTAP features beyond what AWS offers natively.
- Links leverage AWS Lambda to execute code in response to events. This serverless approach removes the dependency of an instance running in your VPC.

#### How links work

Links leverage AWS Lambda. Lambda executes code in response to events and automatically manages the computing resources required by that code. The links that you create are part of your NetApp account and they are associated with an AWS account.

After creating a link, you can associate it with one, or many, FSx for ONTAP file systems. Each file system can be associated only to one link in the same NetApp account. If you have multiple NetApp accounts, a single file system can be associated with additional links under different NetApp accounts.

You create and associate links from the Storage workload in Workload Factory.

You can authenticate links using credentials stored in the Workload Factory credentials service or with your

credentials stored in AWS Secrets Manager. Workload factory doesn't support changing authentication modes.

## Costs

Each transaction that Lambda performs incurs a charge. Because Lambda acts as a proxy between the two systems, there is a charge when Lambda sends a request to the ONTAP REST API on a file system, and when it sends the response back to Workload Factory.

[Learn more about the costs related to using AWS Lambda](#)

## When a link is required

Workload factory requires a link to display some information and to perform some tasks. If you attempt to perform an operation that requires a link and you haven't associated a link with the FSx for ONTAP file system, Workload Factory notifies you that the operation requires a link.

The features that require a link include:

- Well-architected status of FSx for ONTAP file system configurations for proactive maintenance, reliability, and cost-performance optimization
- ONTAP EMS event monitoring and alerting
- NetApp Autonomous Ransomware Protection (ARP/AI)
- Enhanced holistic capacity observability across FSx for ONTAP file systems
- Volume and storage VM data replication, management, and monitoring
- SMB/CIFS shares and NFS export policy provisioning and management
- Management of iSCSI volumes on an FSx for ONTAP file system
- Creation and management of snapshot policies for custom protection SLA
- Inode management enhancements for automatic capacity management
- Volume autogrow for elastic scaling
- Clone creation and management, for instant, in-place, data cloning
- Displaying additional metrics directly from ONTAP such as the ONTAP version

Learn how to [connect a link to an FSx for ONTAP file system](#).

## Connect to an FSx for ONTAP file system with a Lambda link

To perform advanced ONTAP management operations, set up a connection between your Workload Factory account and one or more FSx for ONTAP file systems. This involves associating new and existing Lambda links, and authenticating the links. Link association lets you monitor and manage certain features directly from the FSx for ONTAP file system that are unavailable through the Amazon FSx for ONTAP API.

[Learn more about links](#).

### About this task

Links leverage AWS Lambda to execute code in response to events and automatically manage the computing resources required by that code. The links that you create are part of your NetApp account and they are associated with an AWS account.

You can create a link in your account when defining an FSx for ONTAP file system. The link is used for that file system, and it can be used for other FSx for ONTAP file systems. You can also associate a link for a file system later.

Links require authentication. You can authenticate links using credentials stored in the Workload Factory credentials service or with your credentials stored in AWS Secrets Manager. Only one authentication method is supported per link. For example, if you select link authentication with AWS Secrets Manager, you can't change the authentication method later.



AWS Secrets Manager isn't supported when using a Console agent.

## Associate a new link

Associating a new link includes link creation and association.

You have two options for creating links in this workflow - automatically or manually. You'll need to launch an AWS CloudFormation stack in your AWS account to create the link.

- **Automatically:** Creates a link with automatic registration via Workload Factory. A link created automatically requires tokens for Workload Factory automation and the CloudFormation code is short-lived. It can only be used for up to six hours.
- **Manually:** Creates a link with manual registration using either CloudFormation or Terraform from the Codebox. The code persists giving you more time to complete the operation. This is useful when working with different teams like Security and DevOps that might first need to grant the permissions necessary to complete link creation.

## Before you begin

- You should consider which link creation option you'll use.
- You need to have at least one FSx for ONTAP file system in Workload Factory. To discover FSx for ONTAP file systems, you must have an AWS account with permissions for FSx for ONTAP instances and [add credentials in Workload Factory](#) with *view, planning, and analysis* permissions for Storage management.
- The following ports must be open in the security group associated with the FSx for ONTAP file system for link connectivity.
  - For the Workload Factory console: port 443 (HTTPS)
  - For CloudShell and FSx for ONTAP Emergency Management System (EMS) events analysis: port 22 (SSH)
- The link must be able to connect to the following endpoint: <https://api.workloads.netapp.com>. The web-based console contacts this endpoint to interact with the Workload Factory APIs to manage and operate FSx for ONTAP workloads.
- You must have the following permissions in your AWS account when adding a link using a CloudFormation stack:

```
"cloudformation:GetTemplateSummary",  
"cloudformation:CreateStack",  
"cloudformation>DeleteStack",  
"cloudformation:DescribeStacks",  
"cloudformation>ListStacks",  
"cloudformation:DescribeStackEvents",  
"cloudformation>ListStackResources",  
"ec2:DescribeSubnets",  
"ec2:DescribeSecurityGroups",  
"ec2:DescribeVpcs",  
"iam:ListRoles",  
"iam:GetRolePolicy",  
"iam:GetRole",  
"iam>DeleteRolePolicy",  
"iam:CreateRole",  
"iam:DetachRolePolicy",  
"iam:PassRole",  
"iam:PutRolePolicy",  
"iam>DeleteRole",  
"iam:AttachRolePolicy",  
"lambda:AddPermission",  
"lambda:RemovePermission",  
"lambda:InvokeFunction",  
"lambda:GetFunction",  
"lambda:CreateFunction",  
"lambda>DeleteFunction",  
"lambda:TagResource",  
"codestar-connections:GetSyncConfiguration",  
"ecr:BatchGetImage",  
"ecr:GetDownloadUrlForLayer"
```

## Create automatically

Use CloudFormation to automatically create and register the link within Workload Factory.

### 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 associate a link to and then select **Associate link**.
5. In the Associate link dialog, select **Create a new link** and select **Continue**.
6. On the Create Link page, provide the following:
  - a. **Link name:** Enter the name that you want to use for this link. The name must be unique within your account.
  - b. **AWS Secrets Manager:** Optional. Allows Workload Factory to fetch FSx for ONTAP access credentials from your AWS Secrets Manager.

The link deployment stack automatically adds the following default secret manager ARN regex to the Lambda permission policy:

```
arn:aws:secretsmanager:<link_deployment_region>:<link_deployment_account_id>:secret:FSxSecret*.
```

You can either create secrets in alignment with the default permissions or assign your custom permissions for the link policy.

**Configure VPC private endpoint to AWS Secrets Manager** is disabled by default. Selecting this option stores the secret using the VPC private endpoint instead of storing it locally.

- c. **Link permissions:** Select one of the following options for link permissions:
  - **Automatic:** Select this option so that AWS CloudFormation code automatically creates the Lambda permission policy and execution role.
  - **User-provided:** Select this option to assign a specified Lambda execution role and its attached policies to the Lambda link. The following permissions are required for the Lambda permission policy. The `secretsmanager:GetSecretValue` permission is required only if you enabled AWS Secrets Manager.

```
"ec2:CreateNetworkInterface",  
"ec2:DescribeNetworkInterfaces",  
"ec2>DeleteNetworkInterface",  
"ec2:AssignPrivateIpAddresses",  
"ec2:UnassignPrivateIpAddresses",  
"secretsmanager:GetSecretValue"
```

Enter the Lambda execution role ARN in the text box.

- d. **Tags:** Optionally, add any tags that you want to associate with this link so you can more easily categorize your resources. For example, you could add a tag that identifies this link as being used

by FSx for ONTAP file systems.

Workload factory automatically retrieves the AWS account, location, and security group based on the FSx for ONTAP file system.

7. Select **Create**.

The Redirect to CloudFormation dialog appears and explains how to create the link from the AWS CloudFormation service.

8. Select **Continue** to open the AWS Management Console, and then log in to the AWS account for this FSx for ONTAP file system.

9. On the Quick create stack page, under Capabilities, select **I acknowledge that AWS CloudFormation might create IAM resources**.

Note that three permissions are granted to Lambda when you launch the CloudFormation template. Workload factory uses these permissions when using links.

```
"lambda:InvokeFunction",  
"lambda:GetFunction",  
"lambda:UpdateFunctionCode"
```

10. Select **Create stack** and then Select **Continue**.

You can monitor the link creation status on the Events page. This should take no more than 5 minutes.

11. Return to the Workload Factory interface and you'll see that the link is associated with the FSx for ONTAP file system.

### Create manually

You can create a link using two Infrastructure-as-Code (IaC) tools from the Codebox: CloudFormation or Terraform. With this option, you extract the ARN for the link from AWS CloudFormation and report it here. Workload factory manually registers the link for you.

### 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 actionsenu of the file system to associate a link to and then select **Associate link**.
5. In the Associate link dialog, select **Create a new link** and select **Continue**.
6. On the Create Link page, select CloudFormation or Terraform from the Codebox, and then provide the following:
  - a. **Link name:** Enter the name that you want to use for this link. The name must be unique within your account.
  - b. **AWS Secrets Manager:** Optional. Allows Workload Factory to fetch FSx for ONTAP access credentials from your AWS Secrets Manager.

The link deployment stack automatically adds the following default secret manager ARN regex to the Lambda permission policy:

```
arn:aws:secretsmanager:<link_deployment_region>:<link_deployment_account_id>:secret:FSxSecret*.
```

You can either create secrets in alignment with the default permissions or assign your custom permissions for the link policy.

**Configure VPC private endpoint to AWS Secrets Manager** is disabled by default. Selecting this option stores the secret using the VPC private endpoint instead of storing it locally.

c. **Link permissions:** Select one of the following options for link permissions:

- **Automatic:** Select this option so that AWS CloudFormation code automatically creates the Lambda permission policy and execution role.
- **User-provided:** Select this option to assign a specified Lambda execution role and its attached policies to the Lambda link. The following permissions are required for the Lambda permission policy. The `secretsmanager:GetSecretValue` permission is required only if you enabled AWS Secrets Manager.

```
"ec2:CreateNetworkInterface",  
"ec2:DescribeNetworkInterfaces",  
"ec2>DeleteNetworkInterface",  
"ec2:AssignPrivateIpAddresses",  
"ec2:UnassignPrivateIpAddresses",  
"secretsmanager:GetSecretValue"
```

Enter the Lambda execution role ARN in the text box.

- d. **Tags:** Optionally, add any tags that you want to associate with this link so you can more easily categorize your resources. For example, you could add a tag that identifies this link as being used by FSx for ONTAP file systems.
- e. **Link registration:** Select CloudFormation or Terraform for the instructions for how to register the link, and follow the instructions.

Note that three permissions are granted to Lambda when you launch the CloudFormation template. Workload factory uses these permissions when using links.

```
"lambda:InvokeFunction",  
"lambda:GetFunction",  
"lambda:UpdateFunctionCode"
```

After you successfully create the stack, paste the Lambda ARN in the text box.

- f. Workload factory automatically retrieves the AWS account, location, and security group based on the FSx for ONTAP file system.

7. Select **Create**.

You can monitor the link creation status on the Events page. This should take no more than 5

minutes.

8. Return to the Workload Factory interface and you'll see that the link is associated with the FSx for ONTAP file system.

## Result

Workload factory associates the link with the FSx for ONTAP file system. You can perform advanced ONTAP operations.

## Associate an existing link with an FSx for ONTAP file system

After you create a link, associate it with one or more FSx for ONTAP 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 associate a link to and then select **Associate link**.
5. In the Associate link page, select **Associate an existing link**, select the link, and select **Continue**.
6. Select the authentication mode.
  - Workload Factory: enter the password twice.
  - AWS Secrets Manager: enter the secret ARN.

Ensure that the secret ARN contains the following key valid pairs, though the *filesystemID* is optional.

- `filesystemID = FSx_filesystem_id` (optional)
- `user = FSx_user`
- `password = user_password`



Authentication with AWS Secrets Manager requires a user, either the *FSx\_user* that you provide or another user that was created on the FSx for ONTAP file system. The default user is `fsxadmin` if you don't provide a user.

7. Select **Apply**.

## Result

The link is associated with the FSx for ONTAP file system. You can perform advanced ONTAP operations.

## Troubleshoot issues with AWS Secrets Manager link authentication

### Issue

The link lacks permissions to retrieve the secret.

**Resolution:** Add permissions after the link is active. Log in to the AWS console, locate the Lambda link, and edit the attached permission policy.

**Issue**

The secret isn't found.

**Resolution:** Provide the correct secret ARN.

**Issue**

The secret isn't in the right format.

**Resolution:** Go to AWS Secrets Manager and edit the format.

The secret should contain the following key valid pairs:

- filesystemID = FSx\_filesystem\_id
- username = FSx\_user
- password = user\_password

**Issue**

The secret doesn't contain valid ONTAP credentials for file system authentication.

**Resolution:** Provide credentials that can authenticate FSx for ONTAP file systems in AWS Secrets Manager.

## Manage Workload Factory links

Manage links that you've associated with your Workload Factory account. You can view links that are associated with an FSx for ONTAP file system, provide passwords used for link authentication, and remove links from the Workload Factory console.

[Learn more about links](#) or [create and associate a link](#).

### View the links associated with your account


You can view the links that are currently associated with your account.

**Steps**

1. Log in using one of the [console experiences](#).
2. From the Storage menu, select **Administration** and then **Links**.

Existing links appear on the Links page.

3. To view the FSx for ONTAP file systems that are associated with a link, select the **View** button in the Associated resources section.



4. If you need the Amazon Resource Name (ARN) for the link, you can select the *copy* icon next to the ARN field.

### Edit a link

You can't edit a link from the Workload Factory interface. If you need to make a change to a link, you'll need to create a new link and then associate that link to your file system.



You can edit the Lambda network configuration (for example VPC, subnets, and security groups) using the AWS console and the changes will be reflected in links management UI; however, these changes can lead to connectivity issues between Lambda and ONTAP, and are not recommended.

### Authenticate a link

Provide an administrative user password for Workload Factory credentials or an AWS Secrets Manager secret ARN to connect the link to an FSx for ONTAP file system.

AWS Secrets Manager isn't supported when using a Console agent.



Only one authentication method is supported per link. For example, if you select link authentication with AWS Secrets Manager, you can't change the authentication method later.

### 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 associate a link to and then select **Manage**.
5. In the file system overview, select **Authenticate the link**.
6. In the Authenticate link page, select an authenticate mode:
  - Workload Factory: enter the password twice.
  - AWS Secrets Manager: enter the secret ARN.
7. Select **Apply**.

### Result

The link is authenticated, and you can perform advanced ONTAP operations

### Update the password for link authentication

When the administrative password is invalid, update the password to connect the link to the FSx for ONTAP 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 associate a link to and then select **Manage**.
5. In the file system overview, select **Update password**.
6. In the Authenticate link page, enter the new password twice.
7. Select **Apply**.

### Result

The password is updated, and the link is now connected to the FSx for ONTAP file system.

### Remove a link

You can remove a link that you're no longer using in your environment. Any FSx for ONTAP file systems or other resources that were using the link will be unable to use certain functionality after the link is removed.


Note that the link is only deleted from Workload Factory - it is not deleted from your AWS environment. You must delete the Lambda function from your AWS account after removing the link in Workload Factory.

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **Administration** and then **Links**.

Existing links appear on the Links page.

4. From the Links page, select the actions menu of the link to remove and then select **Remove**.



5. If you are sure, select **Remove** again.

Refer to the AWS documentation to [delete the Lambda function](#).

## Discover cache volumes in Workload Factory

Discover and view *cache* volumes that are associated with cache relationships without leaving the NetApp Workload Factory console. Cache relationships are also known as ONTAP FlexCache relationships. Workload Factory discovers existing cache relationships using FlexCache technology, which is NetApp ONTAP's remote caching capability that accelerates data access, reduces WAN latency, bandwidth and costs for read-intensive workloads, especially where clients need to access the same data repeatedly.

[Learn more about replicating data with FlexCache.](#)

### About this task

Link association is required to discover cache relationships.

A cache relationship can exist between volumes on two ONTAP systems such as one FSx for ONTAP file system and one Cloud Volumes ONTAP system. A cache relationship can also exist within a single FSx for ONTAP file system, from volume to volume.

### Before you begin

Consider the following before you begin.

- You must associate a link to discover cache relationships on a file system. [Learn how to associate an existing link or to create and associate a new link](#). After the link associates, return to this operation.
- You must have an existing cache relationship.

### 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 **Cache relationships** tab.

5. Your cache volumes appear in the table. The table includes the following information about each cache volume:

- **Origin cluster:** The origin, or source, cluster of the FlexCache relationship.
- **Origin volume:** The origin, or source, volume of the FlexCache relationship.
- **Origin storage VM:** The origin, or source, storage VM of the volume.
- **Cache volume:** The cache, or target, volume of the FlexCache relationship.
- **Cache storage VM:** The cache, or target, storage VM of the volume.
- **Status:** The status of the FlexCache relationship.
- **Available storage on cache:** The amount of available storage on the cache volume.
- **Cache file system:** The file system of the cache volume.
- **Write modes:** The write mode of the FlexCache relationship.
- **DR cache:** Indicates whether the FlexCache relationship is a disaster recovery (DR) cache.
- **Export policy:** The export policy of the cache volume.

#### Related information

[Manage cache volumes](#)

## Manage volumes

### Create an FSx for ONTAP volume in Workload Factory

After setting up your FSx for ONTAP file system, create FSx for ONTAP volumes in Workload Factory as virtual resources for grouping your data.

#### About this task

FSx for ONTAP volumes group data virtually, determine how data is stored, and determine the type of access to your data. Volumes don't consume file system storage capacity. The data that is stored in a volume primarily consumes SSD storage. Depending on the volume's tiering policy, the data might also consume capacity pool storage. You set a volume's size when you create it, and you can change its size later.

The following protocols might be used for your volumes:

- SMB/CIFS: file storage protocol for Windows operating systems
- NFS: file storage protocol for Unix operating systems
- iSCSI: block storage protocol

S3 endpoints can be attached to an FSx for ONTAP volume. Using an S3 access point, you can access file data residing on SMB/CIFS or NFS volumes via the AWS S3 APIs. This allows you to integrate your existing data with GenAI, ML, and analytics from AWS services that support S3 access points.

#### Details for volume settings

##### Immutable files

This feature, also known as SnapLock, is disabled by default. Enabling immutable files prevents data deletion or overwriting for a set period. Enabling this feature is possible only during volume creation. After the feature is enabled, it cannot be disabled. This is a premium feature for FSx for ONTAP that carries an additional charge.

For more information, refer to [How SnapLock works](#) in Amazon FSx for NetApp ONTAP documentation.

- **Retention modes:** You can select from two retention modes - *Enterprise* or *Compliance*.
  - In *Enterprise* mode, an immutable files, or SnapLock, administrator can delete a file during its retention period.
  - In *Compliance* mode, a WORM file cannot be deleted before its retention period expires. Similarly, the immutable volume cannot be deleted until the retention periods for all files within the volume expire.
- **Retention period:** The retention period has two settings - *retention policy* and *retention periods*. The *retention policy* defines how long to retain files in an immutable WORM state. You can specify your own retention policy or use the default retention policy (unspecified), which is 30 years. The minimum and maximum *retention periods* define the range of time allowed for locking files.



Even after the retention period expires, you can't modify a WORM file. You can only delete it or set a new retention period to turn on WORM protection again.

- **Autocommit:** You'll have the option to enable the autocommit feature. The autocommit feature commits a file to WORM state on a SnapLock volume if the file did not change for the autocommit period duration. The autocommit feature is disabled by default. You must ensure that the files you want to autocommit reside on a SnapLock volume.
- **Privileged delete:** A SnapLock administrator can turn on privileged delete on a SnapLock Enterprise volume to allow a file to be deleted before the file's retention period expires. This feature is disabled by default.
- **Volume append mode:** You can't modify existing data in a WORM-protected file. However, immutable files allows you to maintain protection for existing data using WORM-appendable files. For example, you can generate log files or preserve audio or video streaming data while writing data to them incrementally. [Learn more about volume-append mode](#) in Amazon FSx for NetApp ONTAP documentation.

## Before you begin

Review the following prerequisites before you create a volume:

- You must have an FSx for ONTAP file system in the Workload Factory console.
- You must have a storage VM.
- For protocol access, complete the following:
  - To configure access to the 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.
  - You must configure access for the protocol you select, either SMB/CIFS, NFS, or iSCSI.

## Create a volume

You can create a volume using the following tools available in the Codebox: REST API, CloudFormation, and Terraform. [Learn how to use Codebox for automation](#).



When using Terraform from Codebox, the code you copy or download hides `fsxadmin` and `vsadmin` passwords. You'll need to re-enter the passwords when you run the code.

## 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 you want to create a volume in, and select **Create volume**.
5. On the Create volume page under General details, provide the following details:
  - a. **Volume name**: Enter a name for the volume.
  - b. **Storage VM name**: Optionally, enter a storage VM name.
  - c. **Volume style**: Select **FlexVol** or **FlexGroup** volume.
    - A *FlexVol volume* is a type of logical storage in ONTAP that allows you to manage data. FlexVol volumes can be moved between aggregates or nodes within the same storage virtual machine (SVM) without disrupting client access. The FlexVol volume style is selected by default.
    - A *FlexGroup volume* is a scale-out NAS container that provides high performance along with automatic load distribution and scalability. It consists of several member volumes (constituents), which are underlying FlexVol volumes that share traffic transparently.
  - d. **Volume size**: Enter the volume size and unit.

Optionally, enable volume autogrow. This option is available when you select **File access** as the volume access type.

+

FlexGroup volume size depends on the number of constituents, requiring 100 GiB per constituent.

- e. **Volume autogrow**: Optionally, enable volume autogrow to automatically expand 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. You cannot set the maximum growth size smaller than the current volume size

- f. **Tags**: Optionally, you can add up to 50 tags.

6. Under Access (only for file systems with associated links), provide the following details:
  - a. **Access type**: Select **File access** or **Block access**. Additional fields to configure volume access differ depending on your selection.
    - **File access**: allows multiple authorized users and devices access to the volume using SMB/CIFS, NFS, or dual (SMB/NFS) protocols.

Complete the following fields to set up file access to the volume.
  - b. **NFS export policy**: Provide the following details to provide NFS access:
    - i. **Access control**: Select a **Custom export policy**, **Existing export policy**, or **No access to the volume** from the dropdown menu.
    - ii. **Export policy name**:

If you selected a custom export policy, select an existing policy name from the dropdown menu.

If you selected an existing export policy, enter a new policy name.
    - iii. **Add Export Policy Rule**: Optionally, for a custom export policy, you can add export policy rules to

the policy.

c. **SMB/CIFS share:** Provide the following:

- i. **Name:** Enter the SMB/CIFS share name to provide access.
- ii. **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".

d. **Security style:** For dual-protocol volumes, select either the UNIX or NTFS security style. UNIX is the default security style for dual-protocol volumes. For detailed guidance on user mapping in this context, refer to the AWS blog article "[Enabling multiprotocol workloads with Amazon FSx for NetApp ONTAP](#)".

- **Block access:** allows hosts running critical business applications access to the volume using the iSCSI protocol. Block access is only available when the file system scale-out deployment has six HA pairs or fewer.

Complete the following fields to set up block access to the volume.

i. **iSCSI configuration:** Provide the following details to configure iSCSI for block access to the volume.

- A. Select **Create a new initiator group** or **Map an existing initiator group**.
- B. Select the **Host operating system** from the dropdown menu.
- C. Enter an **Initiator group name** for a new initiator group.
- D. Under Host Initiators, add one or more iSCSI qualified name (IQN) host initiators.

e. **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 inventory table:** When you enable the inventory table on the volume, the system 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.

f. **S3 access point tags:** Optionally, you can add up to 50 tags or remove tags.

7. Under Efficiency and protection, provide the following details:

a. **Storage efficiency:** Enabled by default. Select to disable the feature.

ONTAP achieves storage efficiency using deduplication and compression features. Deduplication eliminates duplicate data blocks. Data compression compresses the data blocks to reduce the amount of physical storage that is required.

b. **Snapshot policy:** Select the snapshot policy to specify the frequency and retention of snapshots.

The following are default policies from AWS. To display existing snapshot policies, you must [associate a link](#).

## default

This policy automatically creates snapshots on the following schedule, with the oldest snapshot copies deleted to make room for newer copies:

- A maximum of six hourly snapshots taken five minutes past the hour.
- A maximum of two daily snapshots taken Monday through Saturday at 10 minutes after midnight.
- A maximum of two weekly snapshots taken every Sunday at 15 minutes after midnight.



Snapshot times are based on the file system's time zone, which defaults to Coordinated Universal Time (UTC). For information about changing the time zone, refer to [Displaying and setting the system time zone](#) in the NetApp Support documentation.

## default-1weekly

This policy works in the same way as the `default` policy, except that it only retains one snapshot from the weekly schedule.

## none

This policy doesn't take any snapshots. You can assign this policy to volumes to prevent automatic snapshots from being taken.

- c. **Tiering policy:** Select the tiering policy for the data stored in the volume.

*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.

- d. **Immutable files:** Enabling the immutable files feature permanently commits files in this volume to an immutable WORM (write-once-read-many) state. Provide the following details:

- i. Select to enable **Immutable files powered by SnapLock**.

- ii. Select the box to agree and proceed.

- iii. Select **Enable**.

- iv. **Retention mode:** Select **Enterprise** or **Compliance** mode.

- v. **Retention period:**

- Select the retention policy:
  - **Unspecified:** Sets the retention policy to 30 years.
  - **Specify period:** Enter the number of seconds, minutes, hours, days, months, or years to set your own retention policy.
- Select the minimum and maximum retention periods:
  - **Minimum:** Enter the number of seconds, minutes, hours, days, months, or years to set the minimum retention period.
  - **Maximum:** Enter the number of seconds, minutes, hours, days, months, or years to set the maximum retention period.

- vi. **Autocommit:** Disable or enable autocommit. If you enable autocommit, set the autocommit period.
- vii. **Privileged delete:** Disable or enable. If you enable privileged delete, a SnapLock administrator can delete a file before its retention period expires.
- viii. **Volume append mode:** Disable or enable. Enables you to add new content to WORM files.
- e. **ARP/AI:** NetApp Autonomous Ransomware Protection with AI (ARP/AI) is enabled by default when a link is associated with the file system. [Learn more about ARP/AI](#). Accept the statement to proceed.

If the feature is unavailable, it is because of one of the following reasons:

- A link is not associated with the file system. [Learn how to associate an existing link or to create and associate a new link](#). After the link associates, return to this operation.
  - Volumes with immutable files, and volumes with iSCSI and NVMe protocols are not supported for ARP/AI.
  - The file system already has an ARP/AI policy.
8. Under Advance configuration, enter the following details. The fields under this section differ based on the volume style you select.
    - a. **Junction path:** Enter the location in the storage VM's namespace where the volume gets mounted. The default junction path is `/<volume-name>`.
    - b. **Aggregate name:** Only for FlexVol volumes. Select the aggregate to host the new volume.
    - c. **Aggregates list:** Only for FlexGroup volumes. Add or remove aggregates. The minimum number of aggregates is one.
    - d. **Number of constituents:** Only for FlexGroup volumes. Enter the number of constituents per aggregate. 100 GiB is required per constituent.
  9. Select **Create**.

#### Related information

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

## Access your FSx for ONTAP file system data

You can access your FSx for ONTAP file systems from on-premises by mounting volumes for NAS clients and mounting iSCSI LUNs for SAN clients.

[Accessing data](#) in Amazon FSx for NetApp ONTAP documentation provides topics about how to access data for your reference.

You can also get the mount point for volumes in NetApp Workload Factory.

### Get mount point for volumes in NetApp Workload Factory

Get the mount point for a volume to mount a share on a CIFS share or NFS client.

#### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.

3. From **FSx for ONTAP**, select the actions menu of the file system with the volume and 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, then **Basic actions**, and then **View mount command**.
6. In the Mount command dialog, select **Copy** to copy the command for either the NFS mount point or CIFS share. You'll enter the copied command in your terminal.
7. Select **Close**.

### Connect to NAS clients

- [Mount a volume on Linux clients](#)
- [Mount a volume on Windows clients](#)
- [Mount a volume on macOS clients](#)

### Connect to SAN clients

- [Mount an iSCSI LUN on Linux clients](#)
- [Mount an iSCSI LUN on Windows clients](#)

## Create block storage resources

### Create an initiator group for a file system in NetApp Workload Factory

Use NetApp Workload Factory to create initiator groups and manage host access to SAN block devices.

#### About this task

Initiator groups, or igroups, connect block devices (LUNs) to the compute resources that are allowed to access them. Unlike NFS or CIFS, where a volume is broadly accessible and user permissions control access, block storage permissions operate at the machine level. Typically, only one system can access a block device at a time.

An igroup acts as a permission layer for block storage. When a server connects to the storage system, it identifies itself using its iSCSI qualified (IQN) host initiator. If that IQN belongs to one or more igroups, then the server gains access to all LUNs associated with those igroups. Both an igroup and an iSCSI host connection are required for iSCSI to function properly.

#### Before you begin

You must associate a link to create igroups. [Learn how to associate an existing link or to create and associate a new link](#). After you associate the link, return to this operation.

#### Steps

1. Log in using one of the [console experiences](#).
2. In the Storage tile, select **Go to Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the **Block devices** tab.
5. Select the resource type **Create initiator group** and then select **Create igroup**.

6. In the **Create initiator group** dialog, do the following:
  - **igroup name**: Enter a name for the initiator group.
  - **igroup description**: (Optional) Enter a description for the initiator group.
  - **Storage VM name**: Select the storage VM for the initiator group.
  - **Block device name**: Select one or more block devices to associate with the initiator group. The block devices listed are those that have not been mapped to a host initiator yet.
  - **Operating system type**: Select Linux, VMware, or Windows for the operating system type.
  - **Host initiators**: Add one or more iSCSI qualified (IQN) host initiators to the initiator group.
7. Select **Create**.

#### Related information

[Manage the block storage for an FSx for ONTAP file system](#)

## Create a block device for a file system in NetApp Workload Factory

Create block devices to support your line of business (LOB) application requirements.

#### About this task

Only FlexVol volumes are supported for block devices in NetApp Workload Factory. You can create block devices using the iSCSI protocol.

The block size must be smaller than the available FlexVol volume size.

#### Before you begin

- You must associate a link to create block devices. [Learn how to associate an existing link or to create and associate a new link](#). After you associate the link, return to this operation.

#### Steps

1. Log in using one of the [console experiences](#).
2. In the Storage tile, select **Go to Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the **Block devices** tab.
5. Select **Create block device**.
6. Under **Volume details**, do the following:
  - a. **Volume name**: Select one of the following options.
    - Create a new volume and enter the name of the volume.
    - Select an existing volume.
  - b. **Aggregate name**: Select the aggregate to host the new volume.
  - c. **Storage VM**: Select a storage VM.
  - d. **Volume style**: The default volume style is **FlexVol**.
  - e. **Volume size**: Enter the size of the volume and select the unit. The maximum size per FlexVol volume is 100 TiB.
  - f. **Volume autogrow**: Optionally, enable volume autogrow to allow the volume to automatically increase in size when it reaches capacity. The maximum growth size is 300 TiB.

g. **Tags:** Optionally, add tags to help organize and categorize your block device.

7. Under **Block device details**, do the following:

a. **Block device name:** Enter a name for the block device.

b. **Block device size:** Enter the size of the block device and select the unit. The block device size must be smaller than the available volume size.

8. Under **Access**, do the following:

a. **iSCSI configuration:** Select one of the following options.

- **Create new initiator group:** Provide the host operating system, initiator group name, and add one or more iSCSI qualified name (IQN) host initiators.

- **Map existing initiator group:** Select an existing initiator group, provide the host operating system, and select one or more iSCSI qualified name (IQN) host initiators.

9. Under **Efficiency and protection**, do the following:

a. **Storage efficiency:** Enabled by default. Select to disable the feature.

ONTAP achieves storage efficiency using deduplication and compression features. Deduplication eliminates duplicate data blocks. Data compression compresses the data blocks to reduce the amount of physical storage that is required.

b. **Snapshot policy:** Select the snapshot policy to specify the frequency and retention of snapshots.

The following are default policies from AWS. To display existing snapshot policies, you must [xref:./associate a link](#).

#### **default**

This policy automatically creates snapshots on the following schedule, with the oldest snapshot copies deleted to make room for newer copies:

- A maximum of six hourly snapshots taken five minutes past the hour.
- A maximum of two daily snapshots taken Monday through Saturday at 10 minutes after midnight.
- A maximum of two weekly snapshots taken every Sunday at 15 minutes after midnight.



Snapshot times are based on the file system's time zone, which defaults to Coordinated Universal Time (UTC). For information about changing the time zone, refer to [Displaying and setting the system time zone](#) in the NetApp Support documentation.

#### **default-1weekly**

This policy works in the same way as the `default` policy, except that it only retains one snapshot from the weekly schedule.

#### **none**

This policy doesn't take any snapshots. You can assign this policy to volumes to prevent automatic snapshots from being taken.

c. **Tiering policy:** Select the tiering policy for the data stored in the volume.

*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.

- d. **ARP/AI**: NetApp Autonomous Ransomware Protection with AI (ARP/AI) is enabled by default when a link is associated with the file system. [Learn more about ARP/AI](#). Accept the statement to proceed.

If the feature is unavailable, it is because of one of the following reasons:

- A link is not associated with the file system. [Learn how to associate an existing link or to create and associate a new link](#). After the link associates, return to this operation.
- Volumes with immutable files and volumes with the NVMe protocol are not supported for ARP/AI.
- The file system already has an ARP/AI policy.

10. Select **Create**.

#### Related information

[Manage the block storage for an FSx for ONTAP file system](#)

## Create a storage VM for an FSx for ONTAP file system

Create a storage VM (SVM) for an FSx for ONTAP file system to access storage and data services virtually for your workloads in NetApp Workload Factory.

#### About this task

Storage VMs are isolated file servers that you can use to access the data from each workload in Workload Factory Storage. Each SVM has its own administrative credentials and endpoints for administering and accessing data.

With SVMs, when you access data in FSx for ONTAP, your clients and workstations mount a volume, CIFS/SMB share, or iSCSI LUN hosted by an SVM using the SVM's endpoint (IP address).

#### Before you begin

Verify the supported number of storage VMs per file system. Refer to [Managing FSx for ONTAP storage virtual machines](#) in AWS documentation for the maximum number of SVMs per file system.

## Create a storage VM

You can create a storage VM from the Workload Factory console. You can also use the following tools available in the Codebox: REST API, CloudFormation, and Terraform. [Learn how to use Codebox for automation](#).



When using Terraform from Codebox, the code you copy or download hides `fsxadmin` and `vsadmin` passwords. You'll need to re-enter the passwords when you run the code.

#### 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 and then select **Manage**.
5. In the file system overview under Storage VMs, select **Create storage VM**.
6. On the Create storage VM page, under Storage VM configuration, provide the following:
  - a. **Name**: Enter a name for the storage VM.
  - b. **Storage VM credentials**: Provide a password for this storage VM's `vsadmin` user or use the file system's `fsxadmin` user credentials.
  - c. **Root volume security style**: Select the root volume security style depending on the type of clients that access your data - UNIX (Linux clients), NTFS (Windows clients), or Mixed.
  - d. **Tags**: Optionally, you can add up to 50 tags.
7. Select **Create**.

## 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.

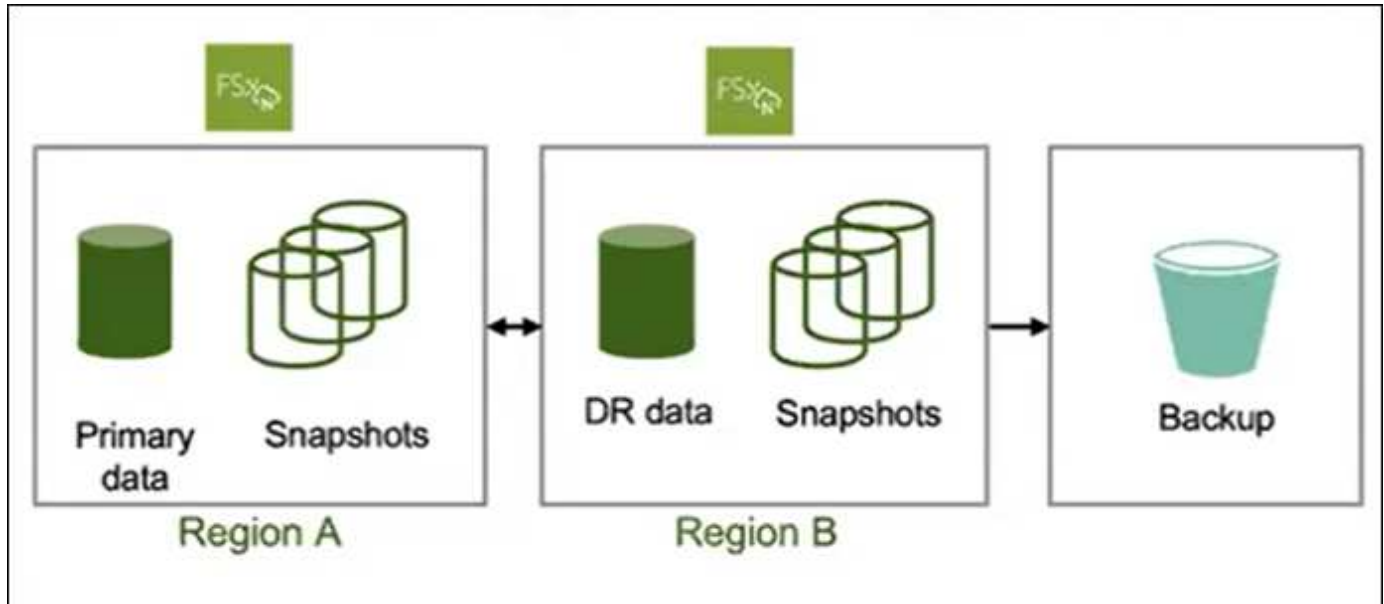
You can choose to create both replicated volumes on another FSx for ONTAP file system and backup files in the cloud. Or you can choose just to create replicated volumes or backup files - it's your choice.

#### 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 volume 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

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

Volume 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, Volume backups**, 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, Volume backups**, 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

### Replicate data to FSx for ONTAP 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. You can replicate data between two FSx for ONTAP file systems, or between an on-premises ONTAP system and an FSx for ONTAP file system.

For storage VM migration, you must complete the cut over operation right after you create a replication relationship.

#### About this task

Replication protects your data if a disaster affects your region; it can also be used for migration purposes.

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

If you replicate a source volume with immutable files, the target volume and file system remain locked until the source volume's retention period 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](#).

### Migration use cases

When you select the migration use case, you can optionally select to replicate storage VM data and configuration settings for a single storage VM. When migrating data and configuration settings simultaneously, ensure that the last replication for the volume completed in the last 24 hours. All volumes in the same storage VM must be selected to use this feature. The tiering policy for all volumes defaults to the tiering policy of the source volume, which is recommended for migration use cases.

Workload Factory supports migration replication between the following storage systems.

- On-premises ONTAP systems and FSx for ONTAP file systems
- Cloud Volumes ONTAP and FSx for ONTAP file systems
- FSx for ONTAP and FSx for ONTAP file systems
  - First to first generation
  - First to second generation
  - Second to second generation

To migrate storage VM data and configuration settings, you must complete two operations.

1. [Create a replication relationship](#), select **Migration** as the use case and select **Replicate storage VM configuration**.
2. [Cut over replication for migration use cases](#) to permanently migrate data and configuration settings from

the source file system to the target FSx for ONTAP file system.

### Create a replication relationship

Replicate data between two FSx for ONTAP file systems, or between an on-premises ONTAP system and an FSx for ONTAP file system.

### Before you begin

Review these requirements 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.
- Replication isn't supported for volumes in a state other than available, created, or misconfigured, and when the ONTAP version isn't compatible.
- For migration use cases ensure that the last replication for the volume completed in the last 24 hours before you create a replication relationship with storage VM data and configuration settings.

### 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 that contains the volume(s) to replicate.
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:
  - a. **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.
    - Artificial Intelligence (AI) and machine learning (ML): replicate your on-premises ONTAP data to FSx for ONTAP and create S3 access points to support AI and ML training in the cloud
    - Migration: transfers your data to the target FSx for ONTAP file system

**Replicate storage VM configuration:** Optionally, select to replicate storage VM data and configuration settings for a single storage VM. When migrating data and configuration settings simultaneously, ensure that the last replication for the volume completed in the last 24 hours. All volumes in the same storage VM must be selected to use this feature. The tiering policy for all volumes defaults to the tiering policy of the source volume, which is recommended for migration use cases.

- 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 *MirrorAllSnapshots*. *MirrorAllSnapshots* 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*.

- b. **FSx for ONTAP file system:** Select credentials, region, and FSx for ONTAP file system name for the target FSx for ONTAP file system.
- c. **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.
- d. **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.
- e. **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:

- 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. 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**.
- c. **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.
- d. **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 inventory table:** When you enable the inventory table on the volume, the system 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.
- e. **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.

If you created a replication relationship for migration purposes, you must cut over all the volumes and their associated storage VM to complete migration of storage VM data and configuration settings to the target FSx for ONTAP file system.

## Cut over replication for migration use cases

After creating a replication relationship for a migration use case, you must cut over replication to complete migration of storage VM data and configuration settings to a target FSx for ONTAP file system. Cutover replication permanently migrates data and storage VM configuration settings from the source file system to the target FSx for ONTAP file system. During the cutover, data gets replicated for the last time. The system deletes the source volume(s) after cutover completes. You can't undo this action.

## Before you begin

Review these requirements before you begin.

- Stop any client access to your storage VM before you cut over replication.
- Ensure all source volumes are not serving any data before you cut over replication.
- Ensure the data is synced between the source and target volumes before you cut over replication.
- 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.

### 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 that contains the volume(s) to replicate.
5. Select the **Replication relationships** tab.
6. In the Replication relationships table, select the replication relationship to cut over and then select **Cut over replication**.
7. Review the information in the Cut over replication dialog and then type *cut over* to confirm.
8. Select **Cut over**.

### Result

After cutover, the source storage VM goes offline.

### Related information

[Modify the tiering policy](#) for the target volume(s) after cutover.

## 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](mailto: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](#).

# Administer and monitor

## Monitor storage operations with Tracker in NetApp Workload Factory

Monitor and track the execution of FSx for ONTAP, credentials, and link operations and monitor task progress with Tracker in NetApp Workload Factory.

### About this task

Workload factory provides Tracker, a monitoring feature, so you can monitor and track the progress and status of FSx for ONTAP, credentials, and link operations, review details for operation tasks and subtasks, and diagnose any issues or failures.

Several actions are available in Tracker. You can filter jobs by time frame (last 24 hours, 7 days, 14 days, or 30 days), workload, status, and user; find jobs using the search function; and download the jobs table as a CSV file. You can refresh Tracker at any time. And you can quickly retry a failed operation or edit parameters for a failed operation and try the operation again.

Tracker supports two levels of monitoring depending on the operation. Each task, such as file system deployment, displays the task description, status, start time, task duration, user, region, proxy resource, task ID, and all related sub tasks. You can view API responses to understand what happened during the operation.

### Tracker task levels with examples

- Level 1 (task): Tracks file system deployment.
- Level 2 (sub task): Tracks the sub tasks related to the file system deployment.

### Operation status

Operation status in Tracker is as follows *in progress*, *success*, and *failed*.

### Operation frequency

Operation frequency is based on the job type and the job schedule.

### Events retention


Events are retained in the user interface for 30 days.

## Track and monitor operations

Track and monitor operations in the NetApp Console with Tracker.

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **Administration** and then **Links**.



4. In the Tracker tab, use the filters or search to narrow job results. You can also download a jobs report.

## View API request

View the API request in the Codebox for a task in Tracker.

### Steps

1. In Tracker, select a task.
2. Select the actions menu and then select **View API request**.

## Retry a failed operation

Retry a failed operation in Tracker. You can also copy the error message of a failed operation.



Only x number of retries are allowed for a failed operation.

### Steps

1. In Tracker, select a failed operation.
2. Select the actions menu and then select **Retry**.

### Result

The operation is re-initiated.

## Edit and retry a failed operation

Edit the parameters of the failed operation and retry the operation outside Tracker.

### Steps

1. In Tracker, select a failed operation.

2. Select the actions menu and then select **Edit and retry**.

You are redirected to the operation page where you can edit the parameters and retry the operation.

## Result

The operation is re-initiated. Go to Tracker to view the status of the operation.

# Implement file system best practices

## Configuration analysis for FSx for ONTAP file systems

NetApp Workload Factory analyzes Amazon FSx for NetApp ONTAP file system configurations regularly to determine if any there are any configuration issues. When issues are found, Workload Factory shows you what the issues are and explains what needs to change to ensure your file system storage achieves peak performance, cost efficiency, and compliance with best practices.

Key capabilities include:

- Daily configuration analysis
- Automatic best practice validations
- Proactive observability
- Insights to action
- AWS Well-Architected Framework advisor

## How it works

Workload Factory analyzes your workloads running on Amazon FSx for NetApp ONTAP file systems deployments daily. The analysis provides well-architected status, insights, and recommendations.

After the daily analysis completes, configurations appear as "optimized" or "not optimized" in the Well-architected dashboard for the deployment. You'll find the total optimization score, configuration issues by category, and a list of configuration issues and recommendations. You can review the recommendations for configuration issues. Some issues can be fixed automatically by Workload Factory, while others require manual intervention. In this case, Workload Factory provides detailed instructions to help you implement the recommended changes.

You can dismiss the analysis of configurations that do not apply to your environments. This avoids unnecessary alerts and inaccurate optimization results. When you dismiss a specific configuration analysis, Workload Factory does not include the configuration in the total optimization score.

## Why it matters

Workload Factory applies best practices to large storage, database, and VMware environments by combining ongoing assessment with recommendation insights and remediation. Automated fixes applied in the Workload Factory console reduce human error, ensure uniform management, and preserve performance and reliability across your workload infrastructures.

## Analysis requirements

For a complete file system analysis, you must do the following:

- Associate a link. Link connectivity lets Workload Factory analyze all file system configurations like data protection and performance.

[Learn how to associate an existing link or to create and associate a new link.](#)

- Grant *view, planning, and analysis* permissions in your AWS account.

[Learn how to grant permissions to an AWS account](#)

## Best practices and recommendations for storage workloads

Workload Factory assesses storage configurations to provide an in-depth view of ONTAP configuration best practices and for compliance with the AWS Well-Architected Framework. The assessment also recommends improvements and fixes.

The well-architected analysis categorizes configurations in the following pillars of the framework: *reliability, security, operational excellence, cost optimization, and performance efficiency*.

### Reliability

Reliability ensures that workloads perform their intended functions correctly and consistently, even when there are disruptions.

- **Schedule volume backups**

Backing up your volumes helps support data retention and compliance needs. Use volume backups to set up automated backups and retention for your data.

- **Schedule local snapshots**

Schedule local snapshots for efficient backup and quick restores. Snapshots are instant, point-in-time images of your volumes.

- **Cross-region replication**

Cross-region replication ensures that your data is replicated to another AWS region, providing enhanced data durability and availability. Workload Factory recommends setting up cross-region replication to help with disaster recovery and compliance.

- **Set up data replication**

To extend data reliability, data can be replicated to an FSx for ONTAP file system in the same region or in another region. Set up data replication to support migration, disaster recovery, and long-term retention across file systems.

- **Increase SSD capacity threshold**

The SSD storage tier capacity should not exceed 80% utilization on an ongoing basis. This might impact data reads and writes to your capacity pool storage tier and impact the throughput capacity of your file system. Running out of capacity might result in data volumes becoming read-only, and services trying to write new data might fail.

- **Match labels to ensure data reliability**

The snapshot policy labels of the source volume and the replication policy labels must match to ensure data reliability.

- **Increase file capacity threshold**

The file capacity threshold should be raised to avoid hitting the volume capacity limit. Low file capacity (inodes) prevents writing additional data to the volume. Workload Factory recommends staying below 80% utilization of the available file capacity on an ongoing basis. Available file capacity is required to create new files in the volume.

## Security

Security emphasizes protecting data, systems, and assets through risk assessments and mitigation strategies.

- **Enable ARP/AI**

NetApp Autonomous Ransomware Protection with AI (ARP/AI) helps protect your volumes from ransomware threats. Workload Factory recommends enabling ARP/AI for all volumes.

- **Unauthorized access to volumes**

Volumes serving application data using iSCSI should not allow NAS access in parallel. Workload Factory recommends that volumes accessed via the iSCSI protocol should be restricted to any additional protocols.

## Operational excellence

Operational excellence focuses on delivering the most optimal architecture and business value.

- **Enable automatic capacity management**

Automatic capacity management should be enabled to regularly ensure that the SSD tier doesn't exceed the threshold.

- **Volume capacity utilization threshold**

Workload Factory recommends that volume capacity doesn't exceed 80% utilization on an ongoing basis. This might impact data reads and writes to your application. Volume capacity increases can be manual or automatic using the volume autogrow feature.

- **Volume utilization nearing full**

When a volume is nearing full capacity, Workload Factory recommends taking action to increase the volume capacity to avoid potential application disruptions.

- **Cache relationship write mode**

For optimal performance, Workload Factory recommends the cache relationship write mode that best suits your workload. Write-around mode provides better performance for read-heavy workloads with small files, whereas write-back mode provides better performance for write-heavy workloads with large files.

- **Optimize cache volume size**

Workload Factory recommends enabling volume autosize and scrubbing on cache volumes to maintain

optimal size and focus the cache on hot data for peak efficiency.

- **Storage VM logical reporting**

Workload Factory recommends that the default reporting setting is set to logical for a storage VM to provide better visibility into storage usage at the volume level.

## Cost optimization

Cost optimization helps you get the most value for your business while keeping costs low.

- **Optimize TCO by tiering cold data**

Cold data tiering should be enabled to reduce SSD storage tier utilization. Applying a tiering policy to every volume is recommended. FSx for ONTAP scans the data continuously to detect cold data and move it to the capacity storage pool tier without disruption.

- **Enable storage efficiencies**

Storage efficiencies should be enabled - compaction, compression, and deduplication - to optimize storage utilization and reduce the SSD tier cost.

- **Unnecessary snapshot and backup deletion**

Snapshots and backups that are no longer needed should be deleted to reduce costs.

- **Inactive block devices**

After a block device isn't used for seven days, Workload Factory recommends archiving block device data or deleting the unused block device to reduce costs.

- **Underutilized SSD capacity**

An underutilized SSD capacity tier may be incurring unnecessary costs. Workload Factory recommends reducing the SSD tier capacity while maintaining a 20% free space buffer above current used capacity. A decrease operation can take several hours to several days, and other file system operations are not supported during this time.

- **Inactive NAS volumes**

The configuration analysis identifies NAS volumes that are not actively used and recommends deleting or archiving the volumes to reduce costs.

## What's next

[Implement well-architected file system configurations](#)

## Implement well-architected file system configurations

Using configuration analysis insights and recommendations, leverage Workload Factory to implement best practices for your FSx for ONTAP file systems. You can easily review the well-architected status, learn about issues with your configurations, take action to improve the architecture of any systems that aren't optimized for reliability, security, efficiency, performance, and cost.

You can also dismiss the analysis of specific storage configurations that don't apply to your storage environment to avoid unnecessary alerts and inaccurate optimization results.

[Learn about the configuration analysis and well-architected status in Workload Factory.](#)

## About this task

Workload factory analyzes Amazon FSx for NetApp ONTAP file system deployment configurations daily. The daily analysis provides the well-architected status, and insights and recommendations with options to automatically fix configuration issues so that your file system meets best practices.

Link connectivity allows Workload Factory to scan for issues with performance, data protection, and configurations. [Connect to an FSx for ONTAP file system using a link](#) for the most comprehensive analysis of your file system resources.

You have options to review the recommendations for configuration issues with your file systems and fix the issues from the Storage within the Workload Factory console.

Because requirements for storage configurations vary, you can dismiss the analysis of specific configurations that don't apply to your storage environment. This helps you avoid unnecessary alerts and inaccurate optimization results. When a specific configuration analysis is dismissed, the configuration isn't included in the total optimization score.

## What is analyzed

Workload factory analyzes the well-architected status of the following configurations for FSx for ONTAP file systems:

- Reliability: SSD capacity threshold, scheduled local snapshots, schedule volume backups, remote data replication, and data reliability for long-term retention
- Security: NetApp Autonomous Ransomware Protection with AI (ARP/AI) disabled and unauthorized access to volumes
- Operational excellence: automatic capacity management, volume file capacity utilization threshold, volume utilization nearing full, cache relationship write mode, optimize cache volume size, and volume logical capacity reporting
- Cost optimization: storage efficiencies, data tiering, unnecessary snapshot and backup deletion, inactive block devices, underutilized SSD capacity, and inactive NAS volumes

## Before you begin

- You must [grant operations and remediation permissions](#) in your AWS account.
- The remediation process may cause instance downtimes or service interruptions. Make sure you review each recommendation carefully before selecting to fix a configuration issue.
- [Connect to an FSx for ONTAP file system using a link](#) for the most comprehensive analysis of your file system resources.

## Fix a configuration issue

You can fix configuration issues for an FSx for ONTAP file system or for selected volumes in a file system. You can select one or more configurations to fix.

## Steps

1. Log in using one of the [console experiences](#).

2. Select the menu and then select **Storage**.
3. From the Storage menu, select **Well-architected**.
4. Select **View issues** for any configuration. Make sure you review the recommendation carefully.

The recommendation explains best practices and potential pitfalls of unoptimized configurations.

5. Select to **Fix**.

When **View and fix** is an option, select the impacted volumes to fix.

6. Review the summary and action items that appear in the dialog to learn what will happen if you choose to fix the issue. Some operations may cause instance downtimes or service interruptions.
7. Select **Continue** to fix the configuration issue.

### **Result**

The process to fix the issue initiates. Select the account settings menu and then select **Tracker** to view the status of the operation.

### **Dismiss a configuration analysis**

Dismiss to stop a configuration analysis indefinitely for an FSx for ONTAP file system or for selected volumes in a file system. You can restart the analysis when needed.

## Dismiss a configuration analysis 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 **Well-architected**.
4. Select **View issues** for any configuration. Make sure you review the recommendation carefully.

The recommendation explains best practices and potential pitfalls of unoptimized configurations.

5. Under Configurations, identify the configuration that doesn't apply to your environment and then select **Dismiss**.
6. In the Dismiss configuration dialog, select **Dismiss** to stop the analysis for the configuration.

## Dismiss a configuration analysis 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 **Well-architected**.
4. Under Configurations, identify the configuration to dismiss for selected volumes and then select **View and fix**.
5. Identify the volume(s) to dismiss from the configuration analysis.
  - For one volume: select the actions menu and then select **Dismiss volume**.
  - For multiple volumes: select the volumes and then select **Dismiss** next to Bulk action.
6. Select **Dismiss** to stop the analysis for the configuration.
7. In the Dismiss volumes dialog, select **Dismiss** to confirm.

## Result

The configuration analysis stops for the file system or selected volumes.

You can reactivate the analysis at any time. The configuration is no longer included in the total optimization score.

## Reactivate a dismissed configuration analysis

Reactivate a dismissed configuration analysis at any time. You can select one or more configurations to reactivate.

## Reactivate a configuration analysis 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 **Well-architected**.
4. Select **Dismissed configurations**.
5. Identify the configuration you want to reactivate and select **Reactivate**.

## Reactivate a configuration analysis 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 **Well-architected**.
4. Select **Dismissed configurations**.
5. Identify the volume(s) to reactivate from the configuration analysis.
  - For one volume: select the actions menu and then select **Reactivate volume**.
  - For multiple volumes: select the volumes and then select **Reactivate** next to Bulk action.

### Result

The configuration analysis is reactivated. A new analysis occurs daily moving forward.

## Analyze FSx for ONTAP EMS events in NetApp Workload Factory

Quickly identify and resolve FSx for ONTAP file system issues with the smart event analyzer in NetApp Workload Factory. The event analyzer automatically extracts and analyzes FSx for ONTAP Event Management System (EMS) events, leveraging Agentic AI with Amazon Bedrock integration.

### About this task

Storage administrators often respond to FSx for ONTAP EMS events only after customer complaints, or by maintaining custom scripts and alarms. This reactive approach can reduce efficiency, delay issue resolution, and increase downtime.

The event analyzer automatically extracts error, alerts, and emergency EMS events from FSx for ONTAP file systems. You can view these events by [connecting to the file system using a link](#) and by [granting view, planning, and analysis permissions](#) in your AWS account. Events are displayed for 72 hours before removal.

With Amazon Bedrock integration, Workload Factory uses AI to analyze events and provide actionable insights to maintain the health and performance of your FSx for ONTAP file systems.

Key benefits include:

- **Advanced troubleshooting:** AI automatically identifies, analyzes, and provides insights to fix FSx for ONTAP EMS events, reducing manual investigation time.
- **Best-practice remediation:** The event analyzer gives clear, actionable steps to resolve FSx for ONTAP EMS events.

When using the event analyzer, you have full control of your environment while benefiting from advanced AI analysis.

To allow Workload Factory to analyze events, you must activate Amazon Bedrock, select the model Workload Factory uses, create a private endpoint to connect to Amazon Bedrock, add permissions, and create an enterprise license.

[Amazon Bedrock pricing](#)

## Data privacy and security

Your data privacy and security are protected through:

- **Data sovereignty:** All data and aggregations stay within your AWS account and are communicated via private VPC endpoint (Amazon Bedrock), with no public internet exposure.
- **No AI Training:** Customer data is not used to train or improve models. Amazon Bedrock processes events in real time but does not train on your data. Results are stored only in your environment.

For more details, refer to the [Amazon Bedrock data protection documentation](#).

## Before you begin

To use the event analyzer, ensure the following:

- You have [operations and remediation permissions](#) in your AWS account to analyze events for FSx for ONTAP file systems.
- Port 22 (SSH) is open in the security group associated with your FSx for ONTAP file system.

Additional requirements (the system will prompt you during log error analysis):

- **Amazon Bedrock model**

Configure Amazon Bedrock APIs for each AWS account. Amazon BedRock APIs are used to provide insights for FSx for ONTAP events.

Recommended model: `anthropic.claude-sonnet-4-20250514-v1:0`. Provide the Inference profile ARN for your selected region.

- **Workload Factory link**

Create and associate a link with an FSx for ONTAP file system to enable AI-powered events analysis. A link establishes a trust relationship between Workload Factory and one or more FSx for ONTAP file systems and leverages AWS Lambda.

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

- **AWS IAM permissions**

Add the following permissions to the Workload Factory IAM role associated policy.

- `bedrock:InvokeModel`
- `bedrock:InvokeModelWithResponseStream`

These permissions allow Workload Factory to invoke Bedrock models for error investigation and remediation guidance. This profile also ensures secure AI access for tailored insights.

Also add the following permissions for AWS credentials associated with Workload Factory:

- `bedrock:GetInferenceProfile`
- `bedrock:ListInferenceProfiles`

These permissions verify model availability.

## View and analyze EMS events for FSx for ONTAP

Use the Workload Factory console to view and analyze EMS events for FSx for ONTAP file systems.

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **Analysis**.
4. From the Analysis screen, select the AWS accounts, credentials, and regions that contain the FSx for ONTAP file systems you want to analyze.

Only FSx for ONTAP file systems with events display on the screen.

5. If needed, complete the AI analysis setup requirements by following the on-screen prompts to meet any missing prerequisites.
6. Find the FSx for ONTAP file system you want to analyze and then select **View events**.
7. Review the detailed event information.

## Volume administration

### Enable volume autogrow in Workload Factory

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

Optionally, you can manually increase the volume capacity of a volume at any time using the [adjust volume capacity feature](#).



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 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 to update and then select **Manage**.
5. From the file system overview, select the **Volumes** tab.
6. From the Volumes tab, select the actions menu for the volume you want to modify.
7. Select **Basic actions**, then **Set volume autogrow**.
8. 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.

9. Select **Apply**.

## Adjust volume capacity in NetApp Workload Factory

Manually adjust the volume capacity of a volume at any time from the NetApp Workload Factory console.

Optionally, you can [enable the autogrow feature](#) to let Workload Factory manage volume capacity for you.

### About this task

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 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 to update and then select **Manage**.
5. From the file system overview, select the **Volumes** tab.
6. From the Volumes tab, select the actions menu of the volume to adjust capacity for.
7. Select **Basic actions**, then **Adjust volume capacity**.
8. In the Adjust volume capacity dialog, set the **Provisioned capacity** and unit.
9. Select **Adjust** to apply the changes.

#### Related information

- [Enable volume autogrow in Workload Factory](#)
- [Rebalance a volume in Workload Factory](#)

## Check and rebalance volume capacity

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.

#### About this task

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 [associate a link](#). If you don't have an existing link, [create a link](#). To associate a link in the file system, select **Associate link** under **Account name**. 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 volumes to rebalance and then select **Manage**.
5. From the file system overview, select the **Volumes** tab.
6. From the Volumes tab, select **Check volume balance** at the top of the table.
7. 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 [associate a link](#). If you don't have an existing link, [create a link](#). To associate a link in the file system, select **Associate link** under **Account name**. 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 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 rebalance and then select **Manage**.
5. From the file system overview, select the **Volumes** tab.
6. From the Volumes tab, select **Check volume balance** at the top of the table.
7. In the Volumes balance window, optionally select **Data distribution** in the FlexVol balance summary to view used capacity per aggregate.
8. Select **Rebalance** to rebalance one or more unbalanced volumes.
9. 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.

10. 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 using one of the [console experiences](#).
2. In **Storage**, select **Go to 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 rebalance and then select **Manage**.
5. From the file system overview, select the **Volumes** tab.

6. From the Volumes tab, select **Check FlexGroup balance** at the top of the table.
7. In the FlexGroup balance window, select **Rebalance** to rebalance one or more unbalanced volumes.
8. 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.
9. **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.

10. Select a layout comparison view and then select **Next**.
  - Volume layout comparison
  - FSx for ONTAP layout comparison
11. Optionally, download a list of volume moves before rebalancing.
12. 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.

## 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 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 to update and then select **Manage**.
5. In the Volumes tab, select the actions menu of the volume to clone.
6. Select **Data protection actions**, then **Manage immutable files**.
7. 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.
8. Click **Apply**.

## Result

The updates now apply to the volume.

## Manage volume tags in NetApp Workload Factory

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 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 to update and then select **Manage**.
5. From the file system overview, select the **Volumes** tab.
6. From the Volumes tab, select the actions menu for the volume to modify tags for.
7. Select **Basic actions** then **Edit volume tags**.

8. On the Edit volume tags page, add, edit, or remove tags.

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

9. Select **Apply**.

## 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 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 cache volume and then select **Manage**.
5. From the file system overview, select the **Cache relationships** tab.
6. Select the actions menu for the cache volume you want to modify and then select **Edit cache name**.
7. 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 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 cache volume and then select **Manage**.

5. From the file system overview, select the **Cache relationships** tab.
6. Select the actions menu for the cache volume you want to modify and then select **Adjust cache capacity**.
7. 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 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 cache volume and then select **Manage**.
5. From the file system overview, select the **Cache relationships** tab.
6. Select the actions menu for the cache volume you want to modify and then select **Edit export policy**.
7. In the **Edit export policy** dialog, change the mount path or select a different export policy to assign to the cache volume.
8. 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 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 cache volume and then select **Manage**.
5. From the file system overview, select the **Cache relationships** tab.
6. Select the actions menu for the cache volume you want to modify and then select **Change caching method**.
7. 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 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 cache volume and then select **Manage**.
5. From the file system overview, select the **Cache relationships** tab.
6. Select the actions menu for the cache volume you want to modify and then select **Prepopulate cache**.
7. 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 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 cache volume and then select **Manage**.
5. From the file system overview, select the **Cache relationships** tab.
6. Select the actions menu for the cache volume you want to delete and then select **Delete cache volume**.
7. 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 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 to update and then select **Manage**.
5. From the file system overview, select the **Volumes** tab.
6. From the Volumes tab, select the actions menu of the volume to change the tiering policy for.
7. Select **Advanced actions**, then **Change tiering policy**.
8. 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**
9. Select **Apply**.

## Update storage efficiency setting of a volume

In NetApp Workload Factory, you can update the storage efficiency setting after volume creation.

### About this task

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 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 to update and then select **Manage**.
5. From the file system overview, select the **Volumes** tab.
6. From the Volumes tab, select the actions menu of the volume to change the tiering policy for.
7. Select **Advanced actions**, then **Set storage efficiency**.
8. Choose to enable or disable volume storage efficiency.
9. Select **Apply** to save the change.

## 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 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 to update and then select **Manage**.
5. From the file system overview, select the **Volumes** tab.
6. From the Volumes tab, select the actions menu for the volume to change the NFS export policy for.
7. Select **Advanced actions**, then **Edit NFS export policy**.
8. 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**.
9. Select **Apply**.

## Manage the SMB/CIFS share for a volume in 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 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 to update and then select **Manage**.
5. From the file system overview, select the **Volumes** tab.
6. From the Volumes tab, select the actions menu of the volume to change the SMB share for.
7. Select **Advanced actions**, then **Manage SMB/CIFS shares**.
8. On the Manage SMB/CIFS shares page, select **Create SMB/CIFS share**.
9. 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".

10. 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 using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
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 using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
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 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 with the volume to update.
5. From the file system overview, select the **Volumes** tab.
6. From the Volumes tab, select the actions menu for the volume you want to use, then select **Advanced actions**, and then **Manage S3 access points**.
7. On the Manage S3 access points page, select **Create and attach S3 access point**.
8. 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.

9. 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 S3 access points for a volume in NetApp Workload Factory](#)
- [Delete S3 access points for a volume in NetApp Workload Factory](#)

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

Update S3 access point details and manage tags 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.

### 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 to update, then select **Manage**.
5. From the file system overview, select the **Volumes** tab.
6. From the Volumes tab, 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**.
7. From the **Manage S3 access points** screen, select the actions menu and then select **Edit access point**.
8. 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](#).

## 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	<ul style="list-style-type: none"><li>netapp-metadata-cloudtrail-cw-role-{uuid}</li><li>netapp-metadata-s3table-integration-role-{uuid}</li></ul>	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 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 to update, then select **Manage**.
5. From the file system overview, select the **Volumes** tab.
6. 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**.
7. From the **Manage S3 access points** screen, select the actions menu and then select **Edit access point**.
8. In the **Edit S3 access point** dialog, ensure that the network configuration is set to **Internet**.
9. Follow the instructions in the dialog to set up the infrastructure for the journal table feature.
10. Download the CloudFormation template.
11. 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.
12. Return to the Workload Factory console and back to the volume to manage S3 access points for.
13. Select **View details** from the volume actions menu.
14. In the **Journal table** tab, enter the CloudTrail ARN.
15. 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- $\langle$ uuid $\rangle$  \
  --policy-name  $\langle$ policy-name $\rangle$ 
aws iam delete-role \
  --role-name netapp-metadata-cloudtrail-cw-role- $\langle$ uuid $\rangle$ 
```

## 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 <code>aws_cloudtrail__data</code> was not created in $\{$ bucket $\}$ . Verify <code>s3table</code> permissions.	<code>S3TableIntegration</code> 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

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

### CloudTrail

Permission	Resource	Why
cloudtrail:CreateTrail	arn:aws:cloudtrail:*:*:trail/netapp-metadata-*	Create the trail
cloudtrail:StartLogging	arn:aws:cloudtrail:*:*:trail/netapp-metadata-*	Enable logging
cloudtrail:AddTags	arn:aws:cloudtrail:*:*:trail/netapp-metadata-*	Apply identification tag
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-*	Pass roles to CloudTrail and CloudWatch Logs
	(condition: PassedToService = cloudtrail.amazonaws.com, logs.amazonaws.com)	

## 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
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

### View the details of S3 access points for a volume in NetApp Workload Factory

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 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 with the volume to view details for.
5. From the file system overview, select the **Volumes** tab.
6. From the Volumes tab, 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**.

7. 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.
8. 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.

## Related information

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

## 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 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 to update, then select **Manage**.
5. From the file system overview, select the **Volumes** tab.
6. 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**.
7. Select the actions menu for the S3 access point to delete and then select **Delete**.
8. 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.

## 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 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 clone to split and then select **Manage**.
5. From the file system overview, select the **Volumes** tab.
6. In the Volumes tab, select the actions menu of the volume with the cloned volume to split.
7. Select **Data protection actions**, then **Split cloned volume**.
8. 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 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 to delete and then select **Manage**.
5. From the file system overview, select the **Volumes** tab.
6. From the Volumes tab, select the actions menu of the volume to delete.
7. Select **Basic actions** then **Delete volume**.
8. 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.

## Block storage administration

### Manage block storage for a file system in NetApp Workload Factory

Use NetApp Workload Factory to simplify managing your block storage resources, including igroups and block devices, and control client access for FSx for ONTAP block devices to provide optimal performance and minimize costs.

#### About this task

From the Workload Factory console, you can manage initiator groups (igroups) to control client access to block devices. You can also view block device details, increase capacity, manage client access, archive data for inactive block devices, and delete block devices.

Check the block storage diagram to check for the following:

- Client connections
- Block device status
- Node-client relationship
- Potential incorrect client connections

### Before you begin

- You must associate a link to manage igroups. [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 igroup or block device to view and manage.

### Manage igroups

You can view igroup details, manage client access, and delete igroups.

#### Manage client access for an igroup

Manage client access for an existing igroup at any time.

#### 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 **Block devices** tab.
5. Select the resource type **Initiator groups (igroups)** to view existing igroups.
6. Go to the actions menu for the block device and select **Manage client access**.
7. Review the client access details displayed for the igroup.
8. To make changes to client access, select **Edit client access**.
9. In the **Edit client access** dialog, you can edit the following:
  - **igroup name**
  - **igroup description**
  - **Storage VM name**
  - **Block device name**
  - **Operating system type**
  - **Host initiators**
10. Select **Apply**.

#### Delete an igroup

Delete an igroup when it is no longer needed.

#### 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 **Block devices** tab.
5. Select the resource type **Initiator groups (igroups)** to view existing igroups.
6. Go to the actions menu for the block device and then select **Delete initiator group**.
7. In the Delete initiator group (igroup) dialog, type "delete" to confirm that you want to delete the igroup, and then select **Delete**.

## Manage block devices for an igroup

Block devices, or LUNs (logical unit numbers), are volumes with file systems in a SAN environment, accessible to hosts over a network.

You can manage block devices for FSx for ONTAP file systems that use the iSCSI protocol.

### View block device details

View details for an existing block device at any time.

#### 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 **Block devices** tab.
5. Navigate to the actions menu for block device and then select **View details**.

General details, consumption, access, and protection information for the block device are displayed.

### Increase the capacity of a block device

Increase the capacity of an existing block device at any time.

#### 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 **Block devices** tab.
5. Navigate to the actions menu for the block device and then select **Increase capacity**.
6. Enter the new capacity for the block device and select the unit.
7. Select **Increase** to apply the changes.



After increasing the size of the block device, follow the procedure your host operating system provides to discover the new size and expand the file system.

## Manage client access for a block device

You can manage client access for an existing block device at any time by creating igroups, and adding or removing block devices and host initiators.

### 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 **Block devices** tab.
5. Navigate to the actions menu for the block device and then select **Manage client access**.
6. If no igroup exists, create a new igroup by selecting **Map igroup > Create igroup**, and then do the following:
  - a. **Block device name**: Enter a block device name. You can select multiple block devices to associate with the igroup.
  - b. **Operating system type**: Select the operating system type.
  - c. **igroup name**: Enter an igroup name.
  - d. **igroup description**: Optionally, enter an igroup description.
  - e. **Host initiators**: Enter one or more host initiators. These initiators must follow iSCSI qualified (IQN) format.
  - f. Select **Create**.
7. If an igroup already exists, select **Map igroup > Map existing igroup** to add or remove block devices and host initiators from the igroup and then select **Map**.

## Archive the data of an inactive block device

Block devices that are no longer mapped to a client or unused for seven consecutive days are classified as inactive block devices. You can archive the data of an inactive block device to the capacity pool tier to reclaim SSD capacity.

### 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 **Block devices** tab.
5. Under Orphan devices, select **View and reclaim capacity**.
6. On the Reclaim space for unused block devices screen, select one or more block devices to archive the data and reclaim the capacity.
7. Select **Archive**.

## Delete a block device

Block devices that are not mapped to a client or are unused for seven days are inactive block devices. This operation unmaps and deletes the selected block device. If the host FlexVol volume does not contain any block devices, Workload Factory deletes it.

## 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 **Block devices** tab.
5. Under Orphan devices, select **View and reclaim capacity**.
6. On the Reclaim space for unused block devices screen, select one or more block devices to archive the data and reclaim the capacity.
7. Select **Delete**.

## Related information

- [Create an igroup for an FSx for ONTAP file system](#)
- [Create a block device for an FSx for ONTAP file system](#)

# File system administration

## Manually adjust file system capacity in Workload Factory

Manually adjust the solid-state drive (SSD) storage capacity of an FSx for ONTAP file system to meet the needs of your project-based workloads with varying active working sets.

Increase SSD storage capacity when usage exceeds your specified threshold. Reduce SSD storage capacity when you are not using the working sets to save money.

Alternatively, you can [enable automatic capacity management](#) so Workload Factory manages file system capacity for you.



Decreasing SSD storage capacity is only supported for second-generation file systems.

### About this task

With elastic file system capacity, you can dynamically adjust the capacity of your file systems to match the needs of your workloads.

Changing file system capacity impacts IOPS for your FSx for ONTAP file system.

When you automatically [provision IOPS](#) for a file system, IOPS increases or decreases by 3 IOPS with every 1 GiB increase or decrease in SSD capacity.

When you [provision IOPS](#) manually, you might need to increase your IOPS allocation to support the increased file system capacity.

For SSD storage capacity limits, refer to [Quotas](#) in AWS FSx for NetApp ONTAP documentation.

### Before you begin

To adjust capacity for a file system, you must first [disable automatic capacity management](#).

## Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Dashboard, select **Adjust SSD capacity**.
4. Select the file system you want to adjust capacity for and then select **Continue**.
5. In the Adjust SSD capacity dialog, enter a number for **Provisioned capacity**.
6. Select the unit for the provisioned capacity.
7. Select **Apply**.

## Manage file system capacity and inodes automatically in Workload Factory

Enabling capacity management lets NetApp Workload Factory automatically add incremental storage to an FSx for ONTAP file system as capacity needs change over time. Additionally, enabling this feature removes the need to monitor capacity manually.

### About this task

The system scans the FSx for ONTAP file system every 30 minutes. It checks if incremental storage needs to be added and verifies available volume inodes. The count increases based on automatic capacity management thresholds.

Enabling inode management lets NetApp Workload Factory automatically add incremental inodes to an FSx for ONTAP file system as capacity needs change. This feature removes the need for manual inode monitoring.

Only one account can manage this feature.

The maximum amount of SSD storage capacity for all FSx for ONTAP file systems is 524,288 GiB. To request a quota increase, refer to [Quotas](#) in AWS FSx for NetApp ONTAP documentation.

### Enable automatic capacity management

Enable automatic capacity management to automatically add incremental storage up to the maximum size limit for an FSx for ONTAP file system.

### Before you begin

Consider the following before you begin:

- You must [grant credentials with the view, planning, and analysis permission policy](#) in Workload Factory to complete this task.
- To make sure volume inodes increase along with storage 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.
- You shouldn't enable this feature during data migration because AWS imposes a minimum six-hour cool down period between SSD capacity increases. This restriction might delay adjustments, so plan accordingly.

### 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 automatic capacity management for.
5. Select **Manage**.
6. Under Information, select the pencil icon next to **Capacity management**. The pencil icon appears next to the drop down arrow when you hover the mouse over the **Capacity management** row.
7. In the **Capacity management** dialog, provide the following:
  - a. **Credentials**: Select credentials with *Automate* permissions from the dropdown menu.
  - b. Select the enable button to **Enable automatic capacity management**.

Alternatively, disable the feature. If you need to increase file system capacity, you must first disable automatic capacity management.

- c. **Mode**: Select the mode to manage capacity automatically. The options are incremental and adaptive modes.
  - **Incremental mode**: increases capacity by a fixed amount at thresholds.
    - **Warning threshold**: Set the warning threshold lower than the threshold increase to trigger a notification from the Workload Factory notification service. The default is 70%.

The warning threshold setting is available only if you [enabled the Workload Factory notification service](#).
    - **Threshold increase**: Enter the maximum percentage increase for the FSx for ONTAP file system. The default is 80%.

This is the threshold at which Workload Factory triggers a job to increase the capacity. For example, if the file system reaches 80% of capacity, then Workload Factory increases capacity.
  - **Adaptive mode**: predicts and changes capacity based on historical data.
    - **Analysis period**: Set the analysis period in hours to determine how much historical data to consider for capacity predictions. The default is 120 hours.
    - **Planned buffer time**: Set the buffer time in hours to ensure capacity is increased before it's actually needed. The default is 48 hours.

8. Select **Apply**.

## Result

A file system scan occurs every 30 minutes to determine whether it needs additional capacity.

## Enable capacity notifications

Set up notifications to get alerts when the file system reaches a certain capacity threshold.

## Before you begin

You must [enable the Workload Factory notification service](#) to use this feature.

## 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 automatic capacity management for.
5. Select **Manage**.
6. Under Information, select the pencil icon next to **Capacity management**. The pencil icon appears next to the drop down arrow when you hover the mouse over the **Capacity management** row.
7. In the **Capacity management** dialog, select **Enable capacity notifications** to turn on notifications when capacity reaches the threshold and whenever Workload Factory performs an automatic capacity increase.
8. Set the **Notification threshold** to receive notifications when the file system reaches a certain capacity percentage. The default is 70%.
9. Select **Apply**.

### Enable inode management for a file system

Enable automatic inode management to increase the number of inodes (files) per volume up to the allowable limit.

A scan of the FSx for ONTAP file system occurs every 30 minutes. It verifies available volume inodes and checks if incremental inodes need to be added. The count increases based on automatic inode management thresholds.



**Terraform users:** Terraform has a limitation that requires that all operations are completed within Terraform. Inode management isn't supported in Terraform, but you can enable automatic inode management in the Workload Factory console.

### Before you begin

Consider the following before you begin:

- To manage volume inodes automatically, 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.
- Automatic inode management can be set up with a *warning threshold* that triggers a notification from the Workload Factory notification service. To use this feature, you must [enable the Workload Factory notification service](#) first.
- You must [grant credentials with the view, planning, and analysis permission policy](#) in Workload Factory to complete this task.

### 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 automatic inodes management for.
5. Select **Manage**.
6. Under Information, select the pencil icon next to **Automatic inodes management**. The pencil icon appears next to the drop down arrow when the mouse hovers over the **Automatic inodes management** row.
7. In the **Automatic inodes management** dialog, provide the following:

- a. **Credentials:** Select credentials with *Automate* permissions from the dropdown menu.
- b. Select the enable button to **Enable automatic inodes management**.

Alternatively, disable the feature. If you need to increase the number of inodes, you must first disable automatic inodes management.

- c. **Warning threshold:** Set the warning threshold lower than the threshold increase to trigger a notification from the Workload Factory notification service. The default is 70%.

The warning threshold setting is available only if you [enabled the Workload Factory notification service](#).

- d. **Threshold increase:** Enter the maximum percentage increase for the number of inodes (files) per volume. The default is 80%.

- e. **Incremental increase:** Enter the percentage to increase the number of inodes (files) incrementally. The default is 10%.

8. Select **Apply**.

### Result

A file system scan occurs every 30 minutes to determine if the volumes need additional inodes (files) per volume.

## Manage FSx for ONTAP file system tags in NetApp Workload Factory

Tags can help you categorize your resources. You can add, edit, and remove tags for a file system at any time in NetApp Workload Factory.

### 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 file system tags**.
5. In the **Manage file system tags** dialog, add, edit, or remove tags as needed.

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

6. Select **Apply**.

## Reset the fsxadmin password in NetApp Workload Factory

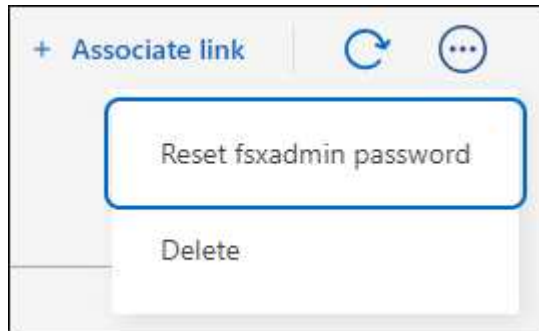
Reset the fsxadmin password in NetApp Workload Factory when necessary.

If you provided an alternate user during file system creation, you can only reset the fsxadmin password in the AWS console.

### 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 reset the fsxadmin password for and then select **Manage**.
5. From the file system overview, select the actions menu.



6. Select **Reset fsxadmin password**.
7. In the Reset fsxadmin password dialog, enter a new fsxadmin password and re-enter it to confirm.
8. Select **Apply**.

## Delete a file system in NetApp Workload Factory

To delete a file system in NetApp Workload Factory, you must first delete any volumes, storage VMs, or replication relationships associated with the 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 FSx for ONTAP file system you want to delete.
5. Select **Manage**.
6. In the **Overview** tab, select the actions menu.
7. Select **Delete**.
8. In the Delete FSx for ONTAP file system dialog, enter the name of the FSx for ONTAP file system to delete.
9. Select **Delete** to confirm.

## Storage VM administration

### Replicate a storage VM to another FSx for ONTAP file system

Replicating a storage VM to another FSx for ONTAP file system in NetApp Workload Factory provides a protective layer of data access in case of data loss. This operation replicates all volumes in one storage VM to another FSx for ONTAP file system.

### Before you begin

To replicate a storage VM to another FSx for ONTAP file system, 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 storage VM to replicate and then select **Manage**.
5. In the file system overview under Storage VMs, select **Manage**.
6. On the Manage storage VMs screen, select the actions menu of the storage VM to replicate an SVM for, then select **Advanced actions > Replicate storage VM**.
7. On the Replicate data 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`.
  - d. **Tiering policy**: Select the tiering policy for the data stored in the target volume.

*Auto* is the default tiering policy when creating a volume using the Workload Factory FSx for ONTAP user interface. 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 MB/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.

8. 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.

- i. For **Choose an existing policy**, select an existing policy from the dropdown menu.
- ii. For **Create a new policy**, provide the following:
  - A. **Policy name**: Enter a policy name.
  - B. **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.

9. Select **Create**.

## Result

All volumes within the storage VM are replicated to the target file system.

## Configure and update Active Directory for a storage VM

Configure and update Active Directory for a storage VM in an FSx for ONTAP file system in NetApp Workload Factory.

### About this task

The same steps apply for configuring and updating Active Directory for a storage VM.

### 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 storage VM to update and then select **Manage**.
5. In the file system overview under Storage VMs, select **Manage**.
6. From the Manage storage VMs screen, select the actions menu of the storage VM to configure Active Directory for, then select **Basic actions > Manage AD configuration**.
7. On the Manage AD configuration page, provide the following:
  - a. **Active Directory domain to join**: Enter the fully qualified domain name (FQDN) of your Active Directory.
  - b. **DNS IP addresses**: Enter up to three IP addresses separated by commas.
  - c. **SMB server NetBIOS name**: Enter the SMB server NetBIOS name of the Active Directory computer object to create for your storage VM. This is the name of this SVM in Active Directory.
  - d. **User name**: Enter the user name of the service account in your existing Active Directory.

Do not include a domain prefix or suffix. For `EXAMPLE\ADMIN`, use `ADMIN`.

- e. **Password**: Enter the password for the service account.
- f. **Organization unit (OU)**: Enter the organization unit.

The OU is the distinguished path name of the organizational unit to which you want to join your file system.

- g. **Delegated administrators group**: Optionally, enter the delegated file system administrators group.

The delegated administrators group is the name of the group in your Active Directory that can administer your file system.

If you are using AWS Managed Microsoft AD, you must specify a group such as AWS Delegated FSx Administrators, AWS Delegated Administrators, or a custom group with delegated permissions to the OU.

If you are connecting to a self-managed AD, use the name of the group in your AD. The default group is `Domain Admins`.

8. Select **Apply**.

## Manage storage VM tags in NetApp Workload Factory

Tags can help you categorize your resources. You can add, edit, and remove tags for a storage VM at any time in NetApp Workload Factory.

### 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 storage VM to update and then select **Manage**.
5. In the file system overview under Storage VMs, select **Manage**.
6. From the Manage storage VMs screen, select the actions menu of the storage VM to edit tags for, then select **Basic actions > Edit storage VM tags**.
7. On the Edit storage VM tags page, add, edit, or remove tags.

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

8. Select **Apply**.

## Reset the storage VM password in NetApp Workload Factory

Reset the password for a storage VM in NetApp Workload Factory when necessary.

### 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 storage VM password to reset and then select **Manage**.
5. In the file system overview under Storage VMs, select **Manage**.
6. From the Manage storage VMs screen, select the actions menu of the storage VM to reset the password for, then select **Basic actions > Reset password**.
7. In the Reset password dialog, provide the following:
  - a. **New password**: Enter a new password for the storage VM.
  - b. **Confirm password**: Enter the new password again to confirm.
8. Select **Apply**.

## Delete a storage VM in NetApp Workload Factory

Delete a storage VM (SVM) that you no longer require from an FSx for ONTAP file system configuration.

### Before you begin

Review the following before you delete a storage VM:

- Make sure that no applications are accessing the data in the SVM.
- Delete all non-root volumes attached to the SVM.

## 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. In the file system overview under Storage VMs, select **Manage**.
6. On the Manage storage VMs screen, select the actions menu of the storage VM to delete.
7. Select **Delete storage VM**.
8. In the Delete storage VM dialog, type “delete” to delete the storage VM.
9. Select **Delete** to confirm.

# Data protection administration

## Snapshots

### Manage snapshot policies

Manage snapshot policies for FSx for ONTAP volumes in Workload Factory. A snapshot policy defines how the system creates snapshots for a volume.

#### About this task

Snapshot management operations like assigning, changing, and deleting snapshot policies for volumes in an FSx for ONTAP file system are managed at the storage VM level. Snapshot policies can be shared with a single storage VM or with all storage VMs.

Some management tasks require you to associate a link with the FSx for ONTAP file system. [Learn about Workload Factory links](#).

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

#### Change a snapshot policy

You can change the snapshot policy name, schedule, and number of copies to retain, and enable or disable immutable snapshots. It isn’t possible to enable or disable policy sharing across storage VMs. This option is available only during snapshot policy creation.

#### Before you begin

To display existing snapshot policies, 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. In the **FSx for ONTAP** tab, select the actions menu of the file system and then select **Manage**.
4. In the file system overview, select the **Storage VMs** tab.
5. From the **Storage VMs** tab, select the actions menu for the storage VM containing the volume to protect with scheduled snapshots, then **Advanced actions**, and then **Manage snapshot policies**.
6. On the Snapshot policy management page, select the actions menu for the snapshot policy to change and then select **Edit**.
7. In the Edit snapshot policy dialog, make the necessary changes to the snapshot policy.
8. Select **Apply**.

## Result

The snapshot policy is updated.

## Enable immutable snapshots

Lock snapshots to prevent them from being deleted during the retention period.

## Before you begin

You must associate a link to enable immutable snapshots. [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. In the **FSx for ONTAP** tab, select the actions menu of the file system that contains the volume to lock snapshots for and then select **Manage**.
4. In the file system overview, select the **Volumes** tab.
5. From the **Volumes** tab, select the actions menu for the volume to protect.
6. Select **Data protection actions**, **Snapshots**, then **Make a snapshot immutable**.
7. In the Make a snapshot immutable dialog, do the following:
  - a. **Snapshot name**: Select the snapshot to lock.
  - b. Set the **Retention period** in number of hours, days, months, or years.
  - c. Accept the statement.
8. Select **Apply**.

## Result

The volume snapshot is now locked.

## Assign a snapshot policy to a volume

You can assign a snapshot policy to a single volume to create scheduled snapshots for the volume.

## Before you begin

You must associate a link to assign a snapshot policy. [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. In the **FSx for ONTAP** tab, select the actions menu of the file system that contains the volume to assign a snapshot policy to and then select **Manage**.
4. In the file system overview, select the **Storage VMs** tab.
5. From the **Storage VMs** tab, select the actions menu for the storage VM containing the volume to protect with scheduled snapshots, then **Advanced actions**, and then **Manage snapshot policies**.
6. On the Snapshot policy management page, select the actions menu of the snapshot policy and then select **Assign policy to volume**.
7. In the Assign snapshot policy dialog, select a snapshot policy to assign to the volume and review the policy schedule.

If the policy contains immutable snapshots, and you want use it, accept the statement.

8. Select **Assign**.

### Result

The snapshot policy is assigned to the volume.

### Remove a snapshot policy from a volume

Remove a snapshot policy from a volume because you no longer want snapshots of the volume or because you want to delete a snapshot policy that is assigned to multiple volumes. To [delete a snapshot policy](#) that is assigned to more than one volume, you must manually remove it from all volumes.

### Before you begin

You must associate a link to remove a snapshot policy. [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. In the **FSx for ONTAP** tab, select the actions menu of the file system that contains the volume to assign a snapshot policy to and then select **Manage**.
4. In the file system overview, select the **Storage VMs** tab.
5. From the **Storage VMs** tab, select the actions menu for the storage VM containing the volume to protect with scheduled snapshots, then **Advanced actions**, and then **Manage snapshot policies**.
6. On the Snapshot policy management page, select the actions menu of the snapshot policy and then select **Assign policy to volume**.
7. In the Assign snapshot policy dialog, select **None** to remove the snapshot policy.
8. Select **Assign**.

### Result

The snapshot policy is removed from the volume.

## Delete a snapshot policy

Delete a snapshot policy when you no longer need it.

When a snapshot policy is assigned to more than one volume, you must manually [remove it](#) from all volumes to delete the snapshot policy. Alternatively, you can [assign a different snapshot policy](#) to the volumes.

### 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 of the storage VM with the snapshot policy to delete, then **Advanced actions**, and then **Manage snapshot policies**.
7. On the Snapshot policy management page, select the actions menu for the snapshot policy to delete and then select **Delete**.
8. In the Delete dialog, select **Delete** to delete the policy.

## Enable and edit snapshots for long-term retention

In NetApp Workload Factory, you can enable snapshots for long-term retention, which lets you replicate specific snapshots for long-term disaster recovery.

Long-term retention enables business services to continue operating even in the event of a complete site failure, supporting transparent failover of applications using a secondary copy.

The same steps apply for enabling and editing snapshots for long-term retention.

When an on-premises ONTAP cluster is the target for the replication relationship, changing snapshots for long-term retention isn't supported.



Editing long-term retention isn't available when replicating storage VM data and configuration settings.

### 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 schedule to change.
7. Select **Edit long-term retention**.
8. In the Edit long-term retention dialog, enable or disable snapshots for long-term retention.
9. If you select to disable snapshots for long-term retention, select **Apply** to complete this operation.

10. If you select to enable snapshots for long-term retention, choose between selecting an existing policy or creating a new policy.
  - a. To use an existing policy, select it from the dropdown menu.
  - b. To create a new policy, provide the following:
    - i. **Policy name**: Enter a policy name.
    - ii. **Snapshot policies**: Select one or more snapshot policies.
    - iii. **Copies to retain**: Enter the number of snapshot copies to retain on the target file system.
11. Select **Apply**.

## Manage snapshots of an FSx for ONTAP volume

Edit snapshot settings, enable directory access, and delete snapshots to manage your snapshots and data protection in Workload Factory.

### Edit a snapshot

Edit the name, label, and retention period of a snapshot. If the snapshot isn't already immutable, you can make the snapshot immutable.

### 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 with the snapshot to edit.
7. Select **Data protection actions** and then **Manage snapshots**.
8. From the Manage snapshots page, select the actions menu for the snapshot to edit, and then select **Edit**.
9. In the Edit a snapshot dialog, you may edit the following:
  - a. Change the name.
  - b. Change the label.
  - c. Change the retention period.
  - d. Optional: **Make this snapshot immutable** to prevent the snapshot from being deleted during the retention period.

If the snapshot is already immutable, you can't edit this setting.

Accept the statement regarding immutable snapshots.

10. Select **Apply**.

### Access a snapshot

Enable snapshot directory access to give users the ability access snapshots autonomously.

### 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 with the snapshot to access.
7. Select **Data protection actions** and then **Manage snapshots**.
8. From the Manage snapshots page, select the actions menu for the snapshot to access, and then select **Access**.
9. In the Access snapshot dialog, select to **Enable snapshot directory access** to access this volume snapshot and all snapshots of the volume.
  - For NFS volumes: Select **NFS access path** to view the NFS path for the snapshot.
  - For SMB/CIFS volumes: Select **SMB access path** to view the SMB path for the snapshot.
10. Copy the access path.
11. Select **Apply**.

#### Restore data from a snapshot

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

#### [Restore a volume from a snapshot](#)

#### Delete a snapshot

Delete a snapshot to free up space.

Immutable snapshots cannot be deleted until the retention period ends.

#### 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 with the snapshot to delete.
7. Select **Data protection actions** and then **Manage snapshots**.
8. From the Manage snapshots page, select the actions menu for the snapshot to delete, and then select **Delete**.
9. In the Delete snapshot dialog, type "delete".
10. Select **Delete** to confirm deletion.

#### Related information

- [Create a snapshot](#)
- [Create a snapshot policy](#)

- [Restore a volume from a snapshot](#)

## Backups

### Manage the backup schedule for an FSx for ONTAP file system

Manage the backup schedule for an FSx for ONTAP file system in NetApp Workload Factory.

#### 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 the backup schedule for and then select **Manage**.
5. Under Information, select the pencil icon next to **Volume backups**. The pencil icon appears next to the dropdown arrow when the mouse hovers over the **Volume backups** row.
6. In the **Volume backups** dialog, provide the following:
  - a. **Daily automatic backups**: Enable or disable the feature. If you disable the feature, select **Apply**. If you enable the feature, complete the following steps.
  - b. **Automatic backup retention period**: Enter the number of days to retain automatic backups.
  - c. **Daily automatic backup window**: Select either **No preference** (a daily backup start time is Selected for you) or **Select start time for daily backups** and specify a start time.
  - d. **Weekly maintenance window**: Select either **No preference** (a weekly maintenance window start time is selected for you) or **Select start time for 30-minute weekly maintenance window** and specify a start time.
7. Select **Apply**.

## Replication

### Replicate data protection volumes in NetApp Workload Factory

Replicate data protection volumes, or cascade the replication of volume data, to extend data protection to tertiary systems or migrate your data.

#### About this task

NetApp Workload Factory supports replicating data protection volumes, also called *cascade deployments*. A *cascade deployment* consists of a chain of relationships in which a source volume is mirrored to a secondary volume (first hop), and the secondary volume is mirrored to a tertiary volume (second hop). If the secondary volume becomes unavailable, you can synchronize the relationship between the primary and tertiary volumes without performing a new baseline transfer.

This feature is supported for FSx for ONTAP file systems with ONTAP version 9.6 and higher. Refer to [ONTAP documentation for compatible ONTAP versions](#).

Learn more about [how cascade deployments work](#).

#### Before you begin

Consider the following before you begin:

- Be aware that volumes that are part of a cascade configuration can take longer to resynchronize.
- If the source volume of the relationship is a data protection volume and is a target of another relationship, reversing the replication relationship isn't supported.
- One replica of a data protection volume (or a second hop) is supported. It isn't considered best practice to create a second replica of a data protection volume (or a third hop).

## 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. From the file system overview, select the **Volumes** tab.
6. In the Volumes table, select one or more data protection volumes (DP/replicated volumes), and then select **Replicate data**.
7. On the Replicate data 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. **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
  - e. **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.

8. 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. 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:** provide the following:
  - **Policy name:** Enter a policy name.
  - Optional: Enable immutable snapshots.
    - 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.

9. Select **Create**.

## Result

The replicated volume or volumes replicate and appear in the **Replication relationships** tab in the target FSx for ONTAP file system.

## Reverse a replication relationship in NetApp Workload Factory

Reverse a replication relationship in NetApp Workload Factory so that the target volume becomes the source volume.

Reverse operations are supported for the following:

- Two FSx for ONTAP file systems
- One FSx for ONTAP file system and one on-premises ONTAP system

After you stop replication and make changes to the target volume, you can replicate those changes back to the

source volume. This process is common in a disaster recovery scenario in which you operate on the target volume for a while and want to switch roles of the volumes.

### About this task

When you reverse and resume a replication, it switches the source and target roles of your volumes; the target volume becomes the new source volume, and the source volume becomes the new target volume. The reverse operation also overwrites the contents of the new target volume with the contents of the new source volume. If you reverse a replication twice, the original replication direction re-establishes.

#### NOTE:

- Any data written to the original source volume between the last data replication and the time that the source volume is disabled is not preserved.
- Reversing replication isn't available when replicating storage VM data and configuration settings.

### Before you begin

Make sure that you know the current and future roles of your source and target volumes because changes on the new target volume are overwritten with the new source volume. If used incorrectly, you can experience unintended data loss.

### 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 reverse.
7. Select **Reverse relationship**.
8. In the Reverse relationship dialog, select **Reverse**.

### Change the replication schedule of a source volume

Change the replication schedule of the source volume in a replication relationship in NetApp Workload Factory.

Choose how frequently snapshots from the source volume are transferred to the replicated volume to match your required point objectives (RPOs).

When an on-premises ONTAP cluster is the target for the replication relationship, changing the replication schedule isn't supported.

### 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 actionsenu of the replication relationship schedule to change.
7. Select **Edit replication interval**.
8. In the Edit replication interval dialog, select the frequency of snapshot transfer from the source volume. You may select between the following frequencies:
  - Every 5 minutes
  - Hourly
  - Every 8 hours
  - Daily
  - Weekly
9. Select **Apply**.

### Limit the max transfer rate of a replication relationship

Limit the max transfer rate of a replication relationship in NetApp Workload Factory. An unlimited transfer rate might negatively impact the performance of other applications and your network.

#### About this task

Limiting the max transfer rate is optional but recommended. 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.

#### Before you begin

Consider how much bandwidth to allocate for replication.

#### 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 actionsenu 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 actionsenu of the replication relationship to limit the max transfer rate for.
7. Select **Edit max transfer rate**.
8. In the Edit max transfer rate dialog, select **Limited** and enter the max transfer limit in MB/s.  
  
Alternatively, select **Unlimited**.
9. Select **Apply**.

## Update snapshot data in a replication relationship

A replication relationship has a set replication schedule, but you can manually update snapshot data transferred between source and target volumes in NetApp Workload Factory at any time.

### 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 update.
7. Select **Update now**.
8. In the Update dialog, select **Update now**.

## Pause and resume a replication relationship in NetApp Workload Factory

Pause a replication relationship to stop scheduled replication updates from the source volume to the target volume. The target volume transitions from read-only to read/write. Both volumes continue to share the latest replication snapshot as a new baseline for later resynchronization.

### About this task

When paused, the replication relationship between source and target volume continues to exist. Data transfers pause and the volumes become independent. To re-enable the transfer of changes from source volume to destination volume, resume the replication.

When you resume a replication, all the changes to the target volume are undone and NetApp Workload Factory re-enables the replication. The target volume transitions from read/write to read-only, and receives updates from the source volume at the scheduled replication interval again. When you resume a replication relationship, the target volume reverts back to the latest initial replication snapshot, at which point, the volume replication process starts over.

### Before you begin

If you pause when a transfer is in progress, the transfer is not affected, and the relationship becomes "Quiescing" until the transfer completes. If the current transfer aborts, it is now a future transfer and will not restart.

## Pause a replication relationship

### 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 pause.
7. Select **Pause (Quiesce)**.
8. In the **Quiesce relationship** dialog, select **Quiesce**.

## Result

The relationship pauses and its status shows as "Paused".

## Resume a paused replication relationship

When you resume a replication relationship, any changes to the destination volume while the replication was stopped are deleted.



Any data written to the original source volume between the last data replication and the time that the source volume is disabled is not preserved.

## 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 resume.
7. Select **Resume**.
8. In the Resume relationship dialog, select **Resume**.

## Result

The relationship resumes and its status shows as "Replicated".

## Stop a replication relationship in NetApp Workload Factory

Stop a replication relationship in NetApp Workload Factory. When you stop a replication relationship, scheduled replication updates from the source volume to the target volume pause. The target volume transitions from read-only to read/write.

## 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 stop.
7. Select **Break**.
8. In the Break replication dialog, select **Break**.

## Result

The replication status of the volume changes to **Broken**. The target volume becomes writable.

## Delete a replication relationship in NetApp Workload Factory

Delete a replication relationship in NetApp Workload Factory. When you delete a replication relationship, it removes the replication relationship between the source and target volume. After the replication relationship deletes, both volumes continue to exist independently with the current data they contain.

When you delete a replication relationship, FSx for ONTAP also deletes the common replication snapshots of the source and target 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 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 delete.
7. Select **Delete**.
8. In the Delete relationship dialog, select **Delete**.

## Performance administration

### Provision SSD IOPS for an FSx for ONTAP file system

Automatically provision or manually provision SSD IOPS for an FSx for ONTAP file system in NetApp Workload Factory.

#### About this task

You can enable automatic SSD IOPS provisioning for an FSx file system or manually provision IOPS.

Automatically provisioned IOPS are calculated as three IOPS per GiB.

If you manually provision IOPS, you might need to increase IOPS before [increasing file system capacity](#).

For information about IOPS limits, refer to [Quotas](#) in AWS FSx for NetApp ONTAP documentation.

### 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 provision IOPS for and then select **Manage**.
5. Under Information, select the pencil icon that appears next to the dropdown arrow when the mouse hovers over the **IOPS allocation** row.

6. In the Provisioned IOPS dialog, select **Automatic** or **User provisioned**.
7. If you select **User provisioned**, enter the desired **IOPS value**.
8. Select **Apply**.

## Update throughput capacity for a file system

Update throughput capacity for an FSx for ONTAP file system in NetApp Workload Factory as needed.

For throughput capacity limits, refer to [Quotas](#) in AWS FSx for NetApp ONTAP documentation.

### 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 throughput capacity for and then select **Manage**.
5. Under Information, select the pencil icon next to **Throughput capacity**. The pencil icon appears next to the drop down arrow when the mouse hovers over the **Throughput capacity** row.
6. In the Throughput capacity dialog, select the throughput capacity you need.
7. Select **Apply** to save the changes.

# Reference

## Performance for FSx for ONTAP in NetApp Workload Factory

For an overview about performance, refer to [Amazon FSx for NetApp ONTAP performance](#) documentation.

## Security for FSx for ONTAP in NetApp Workload Factory

Amazon FSx for NetApp ONTAP documentation provides the following security topics for your reference.

- [Data protection in Amazon FSx for NetApp ONTAP](#)
- [Identity and access management for Amazon FSx for NetApp ONTAP](#)
- [File System Access Control with Amazon VPC](#) in Amazon FSx for NetApp ONTAP documentation

# Knowledge and support

## Register for support

Before you can open a support case with NetApp technical support, you need to add a NetApp Support Site account to Workload Factory and then register for support.

Support registration is required to receive technical support specific to NetApp Workload Factory and its storage solutions and services. You must register for support from the NetApp Console, which is a separate web-based console from Workload Factory.

Registering for support does not enable NetApp support for a cloud provider file service. For technical support related to a cloud provider file service, its infrastructure, or any solution using the service, refer to "Getting help" in the Workload Factory documentation for that product.

[Amazon FSx for ONTAP](#)

## Support registration overview

Registering your account ID support subscription (your 20 digit 960xxxxxxx serial number located on the Support Resources page in the NetApp Console) serves as your single support subscription ID. Each NetApp account-level support subscription must be registered.

Registering enables capabilities like opening support tickets and automatic case generation. Registration is completed by adding NetApp Support Site (NSS) accounts to the NetApp Console as described below.

## Register your account for NetApp support

To register for support and activate support entitlement, one user in your account must associate a NetApp Support Site account with their NetApp Console login. How you register for NetApp support depends on whether you already have a NetApp Support Site (NSS) account.

### Existing customer with an NSS account

If you're a NetApp customer with an NSS account, you simply need to register for support through the NetApp Console.

### Steps

1. In the upper right of the Workload Factory console, select **Help > Support**.

Selecting this option opens the NetApp Console in a new browser tab and loads the Support dashboard.

2. From the NetApp Console menu, select **Administration**, and then select **Credentials**.
3. Select **User Credentials**.
4. Select **Add NSS credentials** and follow the NetApp Support Site (NSS) Authentication prompt.
5. To confirm that the registration process was successful, select the Help icon, and select **Support**.

The **Resources** page should show that your account is registered for support.

 9601111222224444455555 Account Serial Number	 Registered for Support Support Registration
-----------------------------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------------

Note that other NetApp Console users will not see this same support registration status if they have not associated a NetApp Support Site account with their NetApp Console login. However, that doesn't mean that your NetApp account is not registered for support. As long as one user in the account has followed these steps, then your account has been registered.

### Existing customer but no NSS account

If you're an existing NetApp customer with existing licenses and serial numbers but *no* NSS account, you need to create an NSS account and associate it with your NetApp Console login.

#### Steps

1. Create a NetApp Support Site account by completing the [NetApp Support Site User Registration form](#)
  - a. Be sure to select the appropriate User Level, which is typically **NetApp Customer/End User**.
  - b. Be sure to copy the NetApp account serial number (960xxxx) used above for the serial number field. This will speed up the account processing.
2. Associate your new NSS account with your NetApp Console login by completing the steps under [Existing customer with an NSS account](#).

### Brand new to NetApp



If you are brand new to NetApp and you don't have an NSS account, follow each step below.

#### Steps

1. In the upper right of the Workload Factory console, select **Help > Support**.

Selecting this option opens the NetApp Console in a new browser tab and loads the Support dashboard.

2. Locate your account ID serial number from the Support Resources page.

 96015585434285107893 Account serial number	 Not Registered Add your NetApp Support Site (NSS) <a href="#">credentials</a> to BlueXP Follow these <a href="#">instructions</a> to register for support in case you don't have an NSS account yet.
-----------------------------------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

3. Navigate to [NetApp's support registration site](#) and select **I am not a registered NetApp Customer**.
4. Fill out the mandatory fields (those with red asterisks).
5. In the **Product Line** field, select **Cloud Manager** and then select your applicable billing provider.
6. Copy your account serial number from step 2 above, complete the security check, and then confirm that you read NetApp's Global Data Privacy Policy.

An email is immediately sent to the mailbox provided to finalize this secure transaction. Be sure to check your spam folders if the validation email doesn't arrive in few minutes.

7. Confirm the action from within the email.

Confirming submits your request to NetApp and recommends that you create a NetApp Support Site account.

8. Create a NetApp Support Site account by completing the [NetApp Support Site User Registration form](#)
  - a. Be sure to select the appropriate User Level, which is typically **NetApp Customer/End User**.
  - b. Be sure to copy the account serial number (960xxxx) used above for the serial number field. This will speed up the account processing.

#### **After you finish**

NetApp should reach out to you during this process. This is a one-time onboarding exercise for new users.

Once you have your NetApp Support Site account, associate the account with your NetApp Console login by completing the steps under [Existing customer with an NSS account](#).

## **Get help for FSx for ONTAP for Workload Factory**

NetApp provides support for Workload Factory and its cloud services in a variety of ways. Extensive free self-support options are available 24x7, such as knowledgebase (KB) articles and a community forum. Your support registration includes remote technical support via web ticketing.

### **Get support for FSx for ONTAP**

For technical support related to FSx for ONTAP, its infrastructure, or any solution using the service, refer to "Getting help" in the Workload Factory documentation for that product.

#### [Amazon FSx for ONTAP](#)

To receive technical support specific to Workload Factory and its storage solutions and services, use the support options described below.

### **Use self-support options**

These options are available for free, 24 hours a day, 7 days a week:

- [Documentation](#)

The Workload Factory documentation that you're currently viewing.

- [Knowledge base](#)

Search through the Workload Factory knowledge base to find helpful articles to troubleshoot issues.

- [Communities](#)

Join the Workload Factory community to follow ongoing discussions or create new ones.

### **Create a case with NetApp support**

In addition to the self-support options above, you can work with a NetApp Support specialist to resolve any issues after you activate support.

## Before you get started

To use the **Create a Case** capability, you must first register for support. associate your NetApp Support Site credentials with your Workload Factory login. [Learn how to register for support.](#)

## Steps

1. In the upper right of the Workload Factory console, select **Help > Support**.

Selecting this option opens the NetApp Console in a new browser tab and loads the Support dashboard.

2. On the **Resources** page, choose one of the available options under Technical Support:

- a. Select **Call Us** if you'd like to speak with someone on the phone. You'll be directed to a page on netapp.com that lists the phone numbers that you can call.

- b. Select **Create a Case** to open a ticket with a NetApp Support specialist:

- **Service:** Select **Workload Factory**.

- **Case Priority:** Choose the priority for the case, which can be Low, Medium, High, or Critical.


To learn more details about these priorities, hover your mouse over the information icon next to the field name.

- **Issue Description:** Provide a detailed description of your problem, including any applicable error messages or troubleshooting steps that you performed.

- **Additional Email Addresses:** Enter additional email addresses if you'd like to make someone else aware of this issue.

- **Attachment (Optional):** Upload up to five attachments, one at a time.

Attachments are limited to 25 MB per file. The following file extensions are supported: txt, log, pdf, jpg/jpeg, rtf, doc/docx, xls/xlsx, and csv.


ntapitdemo 

NetApp Support Site Account

---

Service Working Enviroment


Select Select

Case Priority 


Low - General guidance



Issue Description

Provide detailed description of problem, applicable error messages and troubleshooting steps taken.

Additional Email Addresses (Optional) 

Type here

Attachment (Optional) Upload 

No files selected  

### After you finish

A pop-up will appear with your support case number. A NetApp Support specialist will review your case and get back to you soon.

For a history of your support cases, you can select **Settings > Timeline** and look for actions named "create support case." A button to the far right lets you expand the action to see details.

It's possible that you might encounter the following error message when trying to create a case:

"You are not authorized to Create a Case against the selected service"

This error could mean that the NSS account and the company of record it's associated with is not the same company of record for the NetApp Console account serial number (ie. 960xxxx) or the system serial number. You can seek assistance using one of the following options:

- Use the in-product chat
- Submit a non-technical case at <https://mysupport.netapp.com/site/help>

## Manage your support cases (Preview)

You can view and manage active and resolved support cases directly from the NetApp Console. You can manage the cases associated with your NSS account and with your company.

Case management is available as a Preview. We plan to refine this experience and add enhancements in upcoming releases. Please send us feedback by using the in-product chat.

Note the following:

- The case management dashboard at the top of the page offers two views:
  - The view on the left shows the total cases opened in the past 3 months by the user NSS account you provided.
  - The view on the right shows the total cases opened in the past 3 months at your company level based on your user NSS account.

The results in the table reflect the cases related to the view that you selected.

- You can add or remove columns of interest and you can filter the contents of columns like Priority and Status. Other columns provide just sorting capabilities.

View the steps below for more details.

- At a per-case level, we offer the ability to update case notes or close a case that is not already in Closed or Pending Closed status.

### Steps

1. In the upper right of the Workload Factory console, select **Help > Support**.

Selecting this option opens the NetApp Console a new browser tab and loads the Support dashboard.

2. Select **Case Management** and if you're prompted, add your NSS account to the NetApp Console.

The **Case management** page shows open cases related to the NSS account that is associated with your NetApp Console user account. This is the same NSS account that appears at the top of the **NSS management** page.

3. Optionally modify the information that displays in the table:

- Under **Organization's cases**, select **View** to view all cases associated with your company.
- Modify the date range by choosing an exact date range or by choosing a different time frame.

Search: Cases opened on the last 3 months Create a case

Date created | Last updated | Last 7 days | Status (5) +  
Last 30 days  
Last 3 months

Date created	Last updated	Priority	Status	Actions
December 22, 2022	December 29, 2022	Medium (P3)	Assigned	...
December 21, 2022	December 28, 2022	Medium (P3)	Active	...
December 15, 2022	December 27, 2022	Medium (P3)	Pending customer	...
December 14, 2022	December 26, 2022	Low (P4)	Solution proposed	...

Apply Reset

- Filter the contents of the columns.

Search: Cases opened on the last 3 months Create a case

Last updated | Priority | Status (5) +

Date created	Priority	Status	Actions
December 29, 2022	Critical (P1)	Active	...
December 28, 2022	High (P2)	Pending customer	...
December 27, 2022	Medium (P3)	Solution proposed	...
December 26, 2022	Low (P4)	Pending closed	...
		Closed	...

Apply Reset

- Change the columns that appear in the table by selecting + and then choosing the columns that you'd like to display.

Search: Cases opened on the last 3 months Create a case

Last updated | Priority | Status (5) +

Date created	Priority	Columns to display	Actions
December 29, 2022	Critical (P1)	Last updated	...
December 28, 2022	High (P2)	Priority	...
December 27, 2022	Medium (P3)	Cluster name	...
December 26, 2022	Low (P4)	Case owner	...
		Opened by	...


Apply Reset

4. Manage an existing case by selecting **⋮** and selecting one of the available options:

- **View case:** View full details about a specific case.
- **Update case notes:** Provide additional details about your problem or select **Upload files** to attach up to a maximum of five files.

Attachments are limited to 25 MB per file. The following file extensions are supported: txt, log, pdf, jpg/jpeg, rtf, doc/docx, xls/xlsx, and csv.

- **Close case:** Provide details about why you're closing the case and select **Close case**.



# Legal notices for NetApp Workload Factory

Legal notices provide access to copyright statements, trademarks, patents, and more.

## Copyright

<https://www.netapp.com/company/legal/copyright/>

## Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<https://www.netapp.com/company/legal/trademarks/>

## Patents

A current list of NetApp owned patents can be found at:

<https://www.netapp.com/pdf.html?item=/media/11887-patentspage.pdf>

## Privacy policy

<https://www.netapp.com/company/legal/privacy-policy/>

## Open source

Notice files provide information about third-party copyright and licenses used in NetApp software.

[NetApp Workload Factory](#)

## 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.