



Amazon FSx for NetApp ONTAP

Workload Factory for NetApp Console

NetApp
October 07, 2026

Table of Contents

- Amazon FSx for NetApp ONTAP 1
 - Use FSx for ONTAP for storage workloads 1
 - Learn about Amazon FSx for NetApp ONTAP in NetApp Workload Factory 1
 - Explore savings with FSx for ONTAP in NetApp Workload Factory 3
 - Create storage resources 7
 - Connect to an FSx for ONTAP file system with a Lambda link 18
 - Discover cache volumes in NetApp Workload Factory 24
 - Access your FSx for ONTAP file system data 26
- Protect your data 26
 - Types of data protection in NetApp Workload Factory 26
 - Use snapshots to protect data 28
 - Use backups to protect data in NetApp Workload Factory 37
 - Use replication to protect data 39
 - Protect your FSx for ONTAP data with NetApp Autonomous Ransomware Protection with AI 54
 - Clone a volume in NetApp Workload Factory 57
 - Protect your FSx for ONTAP data with a cyber vault 57
- Manage FSx for ONTAP in Workload Factory 58
 - Analyze FSx for ONTAP events in NetApp Workload Factory 58
 - Implement file system best practices 61

Amazon FSx for NetApp ONTAP

Use FSx for ONTAP for storage workloads

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.
- **NetApp Terraform 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
- 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](#).

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

With *file system creation* permissions, 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.

With *view, planning, and analysis* permissions, you can copy the content from the Codebox window and deploy the FSx for ONTAP file system using one of the Codebox methods.

Create storage resources

Create or discover cloud storage systems in NetApp Console

Create or discover cloud storage systems to add and manage volumes and additional data services from the NetApp Console.

Create a storage system

The first step is to create a storage system. If you already created a storage system in your cloud provider's console, you can [discover it using the NetApp Console](#).

About this task

A storage VM is created when you create a file system.

Before you begin

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

- The ARN of an IAM role that gives Workload Factory the permissions needed 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.

Steps

1. From the NetApp Console menu, select **Storage** and then **Management**.
2. Select **Add system** from the Systems page.
3. Select **Amazon Web Services** as the location, and then select **Add new** for Amazon FSx for NetApp ONTAP.
Unresolved directive in fsx-storage/create-discover-systems.adoc - include::_include/create-file-system-quick.adoc[]

Result

If you created the file system, the new FSx for ONTAP configuration appears on the Systems page.

You can manage your FSx for ONTAP file systems in several ways, such as from Workloads in the NetApp Console, using ONTAP System Manager, and using AWS CloudFormation.

Advanced create

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

Steps

1. From the NetApp Console menu, select **Storage** and then **Management**.
2. Select **Add system** from the Systems page.
3. Select **Amazon Web Services** as the location, and then select **Add new** for Amazon FSx for NetApp ONTAP.
Unresolved directive in fsx-storage/create-discover-systems.adoc - include::_include/create-file-system-advanced.adoc[]

Result

If you created the file system, the new FSx for ONTAP configuration appears on the Systems page.

You can manage your FSx for ONTAP file systems in several ways, such as from Workloads in the NetApp Console, using ONTAP System Manager, and using AWS CloudFormation.

Discover an existing cloud storage system

If you previously provided your credentials for your cloud storage provider in Workload Factory, you can automatically discover your storage systems from the Discoverable systems page. You can also review available data services.

About this task

You can discover a cloud storage system only once within an account and attach it to one workspace. The storage system can later be removed and re-associated to a different workspace.

Steps

1. From the NetApp Console menu, select **Storage**, then **Management**, and then **Discoverable systems**.
2. The count of discovered cloud storage systems displays. Select **Discover**.
3. Select one or more storage systems and select **Discover** to add them to the Systems page.



- If you select an un-named storage system, you will receive a prompt to enter a name for the system.
- If you select a storage system that doesn't have the credentials required to manage it from the Console, you'll receive a prompt to select credentials with the required permissions.
- The following regions aren't supported for discovery: China regions, GovCloud (US) regions, Secret Cloud, and Top Secret Cloud.

Result

The Console displays your discovered system on the Systems page. You can manage your storage systems in several ways, such as from Workloads in the NetApp Console and using ONTAP System Manager or using your cloud provider's storage management tools.

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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From **FSx for ONTAP**, select the actions menu of the file system and then select **Manage**.
4. In the file system overview under Storage VMs, select **Create storage VM**.
5. 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.
6. Select **Create**.

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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Volumes explorer**.
3. Select **Create volume**.

4. 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.

5. 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.

6. 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:

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

- Select the box to agree and proceed.

- Select **Enable**.

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

- 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.

- Autocommit:** Disable or enable autocommit. If you enable autocommit, set the autocommit period.

- Privileged delete:** Disable or enable. If you enable privileged delete, a SnapLock administrator can delete a file before its retention period expires.

- 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 the NVMe protocol are not supported for ARP/AI. iSCSI volumes are supported beginning with ONTAP 9.17.1.
- The file system already has an ARP/AI policy.

7. Under Advance configuration, enter the following details. The fields under this section differ based on the volume style you select.

- Junction path:** Enter the location in the storage VM's namespace where the volume gets mounted. The default junction path is `/<volume-name>`.
- Aggregate name:** Only for FlexVol volumes. Select the aggregate to host the new volume.
- Aggregates list:** Only for FlexGroup volumes. Add or remove aggregates. The minimum number of aggregates is one.
- Number of constituents:** Only for FlexGroup volumes. Enter the number of constituents per aggregate. 100 GiB is required per constituent.

8. Select **Create**.

Related information

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

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

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

3. From the FSx for ONTAP page, select the **Block devices** tab.
4. Select the resource type **Create initiator group** and then select **Create igroup**.
5. 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.
6. 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. You can create block devices in an existing volume, create a new volume for the new block device, or create a new volume from a snapshot for the new block device.

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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From FSx for ONTAP, select the file system to create a block device for, and then select the **Block devices** tab.
4. Select **Create block device**.
5. 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.

- Create from a volume snapshot. Select the source volume name and snapshot name. Only volumes without block devices are available.
- 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.
6. 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.
 7. 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.
 8. 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 [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.

9. Select **Create**.

Related information

[Manage the block storage for 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 use 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. You use the link for that file system, and you can use it 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.

Link connectivity for GovCloud accounts is supported. Workload Factory uses GovCloud authentication keys, allowing secure, agentless access to your cloud storage systems through GovCloud-approved mechanisms.



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. If you have an AWS GovCloud account, you must create a new link manually.

- **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:

▼ Details

```
[
  "cloudformation:GetTemplateSummary",
  "cloudformation:CreateStack",
  "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:PutRolePolicy",
    "iam>DeleteRole",
    "iam:AttachRolePolicy",
    "lambda:InvokeFunction",
    "lambda:GetFunction",
    "lambda:CreateFunction",
    "lambda>DeleteFunction",
    "iam:PassRole",
    "lambda:TagResource",
    "ecr:BatchGetImage",
    "ecr:GetDownloadUrlForLayer",
    "lambda:AddPermission",
    "lambda:RemovePermission",
    "cloudformation>DeleteStack",
    "kms>CreateGrant",
    "kms:GenerateDataKey"
]

```

Create automatically

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



Automatic link creation isn't supported for AWS GovCloud accounts.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Administration > Links**.
3. On the Links page, select **Create and associate**.
4. In the Associate Link dialog, select the file system to associate with the link and select **Continue**.
5. Select **Create a new link** and then select **Continue**.
6. On the Create Link page, select **Create automatically** 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 execution role:** 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",
```

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**.

The CloudFormation template grants Lambda three permissions when you launch it. 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 five 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 to the [NetApp Console](#).
2. Select the menu , **Administration > Links**.
3. On the Links page, select **Create and associate**.
4. In the Associate Link dialog, select the file system to associate with the link and select **Continue**.
5. Select **Create a new link** and then select **Continue**.
6. On the Create Link page, select **Create manually** and then select CloudFormation or Terraform from the Codebox.
7. Under Details and credentials, 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:** Always enabled for GovCloud accounts. Optional for other accounts. Allows Workload Factory to fetch FSx for ONTAP access credentials from your AWS Secrets Manager.

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 when using AWS Secrets Manager.

```
"ec2:CreateNetworkInterface",  
"ec2:DescribeNetworkInterfaces",
```

```
"ec2:DeleteNetworkInterface",  
"ec2:AssignPrivateIpAddresses",  
"ec2:UnassignPrivateIpAddresses"
```

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.

The CloudFormation template grants Lambda three permissions when you launch it. Workload Factory uses these permissions when using links.

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

- i. After you successfully create the stack, paste the Lambda ARN in the text box in the Workload Factory console. Workload Factory automatically retrieves the AWS account, location, and security group based on the FSx for ONTAP file system.
- ii. Select **Create**.

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

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 systems.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From the FSx for ONTAP page, select the actions menu of the file system to associate a link to and then select **Associate link**.
4. In the Associate link page, select **Associate an existing link**, select the link, and select **Continue**.
5. Select the authentication mode.
 - Workload Factory: enter the password twice.
 - AWS Secrets Manager: enter the secret ARN.

Verify that the secret ARN contains the following key-value 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.

6. Select **Apply**.

Result

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

Related information

[Learn more about links.](#)

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-value 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.

Discover cache volumes in NetApp 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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From **FSx for ONTAP**, select the **Cache relationships** tab.
4. 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](#)

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 to the [NetApp Console](#).
2. Select the menu , **Workloads** > **Volumes explorer** > **FSx for ONTAP**.
3. Select the actions menu for the volume, then **Basic actions**, and then **View mount command**.
4. 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.
5. 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](#)

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 to protect data

Use snapshots to protect data in NetApp Workload Factory

Protect volume data from accidental deletion or corruption by creating and managing

snapshots in NetApp Workload Factory. You can create an on-demand snapshot of a volume, define a snapshot policy to automate snapshot creation and retention for storage VMs, and restore a volume from an existing snapshot to recover data at a specific point in time.

Create an on-demand snapshot of a 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 to the [NetApp Console](#).
2. Select the menu , **Workloads** > **Storage** > **FSx for ONTAP**.
3. From **FSx for ONTAP**, select the actions menu of the file system that contains the volume to create a snapshot 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 with snapshots.
6. Select **Data protection actions** and then **Manage snapshots**.
7. From the Manage snapshots page, select **Create snapshot**.
8. 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.

9. Select **Create**.

Create a snapshot policy for storage VMs

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.

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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From **FSx for ONTAP**, select the actions menu of the file system with the volume 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 volume to protect with scheduled snapshots, then **Advanced actions**, and then **Manage snapshot policies**.
6. On the Snapshot policy management page, select **Create snapshot policy**.
7. In the **Snapshot policy name** field, enter a name for the snapshot policy.
8. Optionally, enter a description for the snapshot policy.
9. Under **Policy schedule and copies**, select when to create snapshots. For example, every minute or hourly.

You can select more than one frequency.

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

The maximum number of copies is 1,023.

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

12. **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.

13. 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.

14. **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.

15. Select **Create**.

Restore a volume from a snapshot

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.

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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From **FSx for ONTAP**, select the actions menu of the file system with the volume 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 restore from a snapshot.
6. Select **Data protection actions** and then **Manage snapshots**.
7. From the Manage snapshots page, select the actions menu for the snapshot to restore, and then select **Restore**.
8. 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.

9. Select **Restore**.

Manage snapshot policies for volumes in NetApp Workload Factory

Manage snapshot policies for FSx for ONTAP volumes in Workload Factory. A snapshot policy defines how the system creates snapshots for a volume.

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.

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 to the [NetApp Console](#).
2. Select the menu , **Workloads** > **Storage**.
3. From the Storage menu, select **Inventory**.
4. Select the actions menu of the file system that contains the volume to assign a snapshot policy to 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 storage VM containing the volume to protect with scheduled snapshots, then **Advanced actions**, and then **Manage snapshot policies**.
7. On the Snapshot policy management page, select the actions menu of the snapshot policy and then select **Assign policy to volume**.
8. 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.

9. Select **Assign**.

Result

The snapshot policy is assigned to the volume.

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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage**.
3. From the Storage menu, select **Inventory**.
4. Select the actions menu of the file system 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 storage VM containing the volume to protect with scheduled snapshots, then **Advanced actions**, and then **Manage snapshot policies**.
7. On the Snapshot policy management page, select the actions menu for the snapshot policy to change and then select **Edit**.
8. In the Edit snapshot policy dialog, make the necessary changes to the snapshot policy.
9. 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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Volumes explorer**.
3. Select the actions menu of the volume to lock snapshots for.
4. Select **Data protection actions**, **Snapshots**, then **Make a snapshot immutable**.
5. 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.
6. Select **Apply**.

Result

The volume snapshot is now locked.

Enable long-term snapshot retention in NetApp Workload Factory

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 to the [NetApp Console](#).
2. Select the menu , **Workloads** > **Storage** > **FSx for ONTAP**.
3. From the **FSx for ONTAP** menu, select **Inventory** and then select the actions menu of the file system to update and then select **Manage**.
4. From the file system overview, select the **Replication relationships** tab.
5. In the Replication relationships tab, select the actions menu of the replication relationship schedule to change.
6. Select **Edit long-term retention**.
7. In the Edit long-term retention dialog, enable or disable snapshots for long-term retention.
8. If you select to disable snapshots for long-term retention, select **Apply** to complete this operation.
9. 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.
10. Select **Apply**.

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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From the **FSx for ONTAP** menu, select **Inventory** and then select the actions menu of the file system to update 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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From the **FSx for ONTAP** menu, select **Inventory** and then select the actions menu of the file system to update 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 of the storage VM with the snapshot policy to delete, then **Advanced actions**, and then **Manage snapshot policies**.
6. On the Snapshot policy management page, select the actions menu for the snapshot policy to delete and then select **Delete**.
7. In the Delete dialog, select **Delete** to delete the policy.

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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Volumes explorer**.
3. Select the actions menu for the volume with the snapshot to edit.
4. Select **Data protection actions** and then **Manage snapshots**.
5. From the Manage snapshots page, select the actions menu for the snapshot to edit, and then select **Edit**.
6. 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.

7. Select **Apply**.

Access a snapshot

Enable snapshot directory access to give users the ability access snapshots autonomously.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Volumes explorer**.
3. Select the actions menu for the volume with the snapshot to access.
4. Select **Data protection actions** and then **Manage snapshots**.
5. From the Manage snapshots page, select the actions menu for the snapshot to access, and then select **Access**.
6. 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.
7. Copy the access path.
8. 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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Volumes explorer**.
3. Select the actions menu for the volume with the snapshot to delete.
4. Select **Data protection actions** and then **Manage snapshots**.
5. From the Manage snapshots page, select the actions menu for the snapshot to delete, and then select **Delete**.
6. In the Delete snapshot dialog, type "delete".
7. Select **Delete** to confirm deletion.

Related information

- [Create a snapshot](#)
- [Create a snapshot policy](#)
- [Restore a volume from a snapshot](#)

Use backups to protect data in NetApp Workload Factory

Protect your volume data from accidental deletion, corruption, and other data loss events by creating and restoring backups in NetApp Workload Factory. You can set up a backup schedule, create an on-demand backup of a volume outside your regularly scheduled backups, and you can restore a volume from an existing backup to any FSx for ONTAP file system in your AWS account.

Create an on-demand backup of a volume

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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Volumes explorer**.
3. Select the actions menu for the volume to back up.
4. Select **Data protection actions**, **Volume backups**, and then **Manual backup**.
5. In the Manual backup dialog, enter a name for the backup.
6. Select **Back up**.

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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From **FSx for ONTAP**, select the actions menu of the file system to update the backup schedule for and then select **Manage**.
4. 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.
5. 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.
6. Select **Apply**.

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

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

3. Select the actions menu for the volume to restore from a backup.
4. Select **Data protection actions**, **Volume backups**, and then **Restore from a backup**.
5. 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.
6. 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.

7. Select **Restore**.

Use replication to protect data

Replicate data to FSx for ONTAP in NetApp Workload Factory

Replicate data in NetApp Workload Factory to protect against data loss if the region where your data resides experiences a disaster. Data replication is supported between FSx for ONTAP file systems and on-premises ONTAP systems or Cloud Volumes ONTAP. Workload Factory also supports cascading replication to extend data protection to tertiary systems or migrate your data.

When replicating storage VM data and configuration settings, you must first create and initialize the replication relationship and then wait for the data to sync at least once before completing the cut over operation.

Standard replication

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

Replication includes creating and initializing a replication relationship. To set up replication, create the replication relationship. To transfer the data between the source and target storage systems, initialize the replication relationship. When migrating storage VM data and configuration settings, you must also complete the cutover operation.

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.



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

For migration, you can choose to replicate storage VM data and configuration settings for one storage VM. If you migrate data and configuration settings together, make sure the last replication for the volume finished within the last 24 hours. You must select all volumes in the same storage VM 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 for storage VMs 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
 - Second to second generation

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

1. [Create the replication relationship](#), select **Migration** as the use case and select **Replicate storage VM configuration**.
2. [Initialize the replication relationship](#) to transfer the data from the source to the target storage system.
3. [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.

Step 1: Create the replication relationship

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

Before you begin

Review these requirements before you begin.

- Ensure that one FSx for ONTAP file system is available to use for the target in the replication relationship.
- Open TCP ports 11104 and 11105 for secure data transfer and management between storage systems. ICMP is optional, but can help with network checks and troubleshooting.
- 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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From the FSx for ONTAP page, select the file system that contains the volume(s) to replicate.
4. 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**.

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

Unresolved directive in fsx-storage/create-replication.adoc - include::_include/create-replication.adoc[]

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.

Step 2: Initialize the replication relationship

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

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

Unresolved directive in fsx-storage/create-replication.adoc - include::_include/initialize-replication.adoc[]

3. In the Initialize relationship dialog, select **Initialize**.

When the replication status shows **Idle**, the data transfer is complete.

Step 3: Cut over replication for migration use cases

After creating the replication relationship for a migration use case, you must cut over replication to complete migration of storage VM data and configuration settings to the 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 source storage VM goes offline when the cutover begins.

Plan for the cutover operation

Consider the following when planning for the cutover operation for migration use cases.

- **Replication version support:** Source and target storage systems must have the same major ONTAP version. For example, if the source system is running ONTAP version 9.18.1, then the target system must also run a version of 9.18.1.
- **Storage VM configuration settings:**
 - The cutover operation doesn't migrate all storage VM configuration settings to the target storage VM. For example, iSCSI, broadcast-domain MTU, and DNS settings do not migrate. Update these settings manually on the target storage VM after cutover. For more details, refer to the [Update configuration settings after cutover](#) section.
 - If your storage VM uses domain-tunnel authentication, set up and enable a domain tunnel on the target file system before you start the migration.
- **Clones:** Clones get migrated as individual volumes, and therefore, will consume more space on the target FSx for ONTAP file system. Consider deleting the clones before cutover to avoid increased capacity on the target.
- **Cache volumes:** For ONTAP FlexCache 9.17.x or below. Source storage VMs that contain cache volumes aren't supported for cutover. Remove the ONTAP FlexCache relationship before cutover to avoid migration failures.
- **AWS S3 access points:** Cutover replication for migration use cases isn't supported for volumes with S3 access points. Delete S3 access points from the source volume to avoid migration failures.
- **ONTAP S3 configurations:** Cutover replication for migration use cases isn't supported for vservers with ONTAP S3 configurations. Remove ONTAP S3 configurations including object-store-server, buckets, users, groups, or policies on the vserver, to avoid migration failures.
- **LUN serial numbers and iSCSI target IQN:** LUN serial numbers and the iSCSI target IQN change during cutover.

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.
 - The following ports must be open in the security group associated with the FSx for ONTAP file system for link connectivity: port 443 (HTTPS) and port 22 (SSH).

Steps

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

Unresolved directive in fsx-storage/create-replication.adoc - include::_include/cutover-replication.adoc[]

8. Select **Cut over**.

Result

After initiating cutover, the source storage VM goes offline.

Update configuration settings after cutover

After you cut over replication for migration use cases, update the following storage VM configuration settings on the target storage VM.

- For iSCSI:
 - Recreate the igroup
 - Update LUN mapping
 - Log in to each SAN client and update iSCSI/MPIO settings for the new target and remove the old vserver target IPs
- Validate broadcast-domain MTU settings are as expected
- Validate DNS settings are as expected
- Recreate cache relationships (ONTAP FlexCache relationships) if there were cache volumes in the source storage VM.

If needed, [modify the tiering policy](#) for the target volume(s) after cutover.

Cascade replication for data protection volumes

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).
- Open TCP ports 11104 and 11105 for secure data transfer and management between storage systems. ICMP is optional, but can help with network checks and troubleshooting.

Steps

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

2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From **FSx for ONTAP**, select the actions menu of the file system that contains the volume(s) to replicate and then select **Manage**.
4. From the file system overview, select the **Volumes** tab.
5. In the Volumes table, select one or more data protection volumes (DP/replicated volumes), and then select **Replicate data**.
6. On the Replicate data page, under Replication target, provide the following:

Unresolved directive in fsx-storage/create-replication.adoc - include::_include/create-replication.adoc[]

1. Select **Create**.

Result

The replicated volume or volumes replicate and appear in the **Replication relationships** tab in the target FSx for ONTAP file system.

Related information

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

Replicate 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 as a backup in the cloud.

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

- **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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From the **FSx for ONTAP** menu, select **On-premises ONTAP**.
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**: Provide the following:
 - GovCloud accounts: Enter ARN with ONTAP credentials.
 - All other accounts: 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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From the **FSx for ONTAP** 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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From the **FSx for ONTAP** 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.

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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From **FSx for ONTAP**, select the actions menu of the file system with the storage VM to replicate and then select **Manage**.
4. In the file system overview under Storage VMs, select **Manage**.
5. 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**.
6. 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.

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.

- 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.

8. Select **Create**.

Result

All volumes within the storage VM are replicated to the target file system.

Track replication delays in NetApp Workload Factory

Turn on tracking to monitor replication delays across your account. Workload Factory alerts you when delays go beyond a set threshold. By default, the threshold is 150% of the replication interval to allow normal transfer time and reduce false alarms. You can adjust the threshold to match your RPO and get notifications when delays exceed it. You can also export the data to a CSV file for further analysis.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From **FSx for ONTAP**, select the actions menu of the file system to update and then select **Manage**.
4. From the file system overview, select the **Replication relationships** tab.
5. Select **Enable tracking**.
6. In the Enable replication lag history dialog, select **Enable**.
7. Back in the Replication lag history tile, select **Customize threshold** to set your preferred threshold.
8. In the Customize lag time threshold dialog, enter the threshold in hours and minutes.
9. Optionally, select **Enable lag time notifications** to receive notifications when replication lags exceed the threshold.

Notifications come in an email and include:

- Source file system
- Source volume
- Target file system
- Target volume
- Lag time threshold that was exceeded
- Link to the replication relationship

10. Select **Apply**.

Related information

[Manage replication relationships](#)

Manage your replication relationships in NetApp Workload Factory

Manage your replication relationships in NetApp Workload Factory. You can pause, resume, reverse, stop, and delete replication relationships between source and target volumes. You can also change the replication schedule of the source volume, and limit the max transfer rate for your replication relationships.

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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage**.
3. From the Storage menu, select **Inventory**.
4. In the **Inventory**, 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 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

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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage**.
3. From the Storage menu, select **Inventory**.
4. In the **Inventory**, 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 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**.

Manually 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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage**.
3. From the Storage menu, select **Inventory**.
4. In the **Inventory**, 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**.
8. In the Update dialog, select **Update**.

Pause a replication relationship

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.

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.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage**.
3. From the Storage menu, select **Inventory**.
4. In the **Inventory**, 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. The target volume transitions from read/write to read-only, and receives updates from the source volume at the scheduled replication interval again.



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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage**.
3. From the Storage menu, select **Inventory**.
4. In the **Inventory**, 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".

Reverse a replication relationship

Reverse a replication relationship 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.

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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage**.
3. From the Storage menu, select **Inventory**.
4. In the **Inventory**, 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**.

Stop a replication relationship

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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage**.
3. From the Storage menu, select **Inventory**.
4. In the **Inventory**, 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

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

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

3. From the Storage menu, select **Inventory**.
4. In the **Inventory**, 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**.

Protect your FSx for ONTAP data with NetApp Autonomous Ransomware Protection with AI

Protect your data with NetApp Autonomous Ransomware Protection with AI (ARP/AI). It monitors activity in NAS (NFS/SMB) and SAN environments to detect unusual behavior that could signal a ransomware attack. If a threat is detected, ARP/AI automatically creates new immutable snapshots so 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 (used in both NAS and SAN environments)
 - **File extension types**: An extension that does not conform to the normal extension type (used in NAS environments only)
 - **File IOPS**: A surge in abnormal volume activity with data encryption (used in NAS environments only)

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 behavior depends on the protocol and volume type:

- **NAS volumes (NFS/SMB)**: ARP/AI starts working as soon as you enable it. It uses a pre-trained AI model to detect ransomware threats immediately.
- **SAN volumes (beginning with ONTAP 9.17.1)**: ARP/AI provides immediate protection, even during the first evaluation period. During this 2-4 week period, the system learns normal encryption activity and sets alert thresholds. At the same time, it continues to monitor your data, detect threats, and send alerts.
- **Active mode**: When active mode is on, FSx for ONTAP creates ARP/AI snapshots to help protect your data when it detects a possible threat.

If the system flags a file extension as unusual, review the alert. You can respond to the alert to protect your data, or mark it as a false positive if the activity is expected.

When you mark an alert as a false positive, the system updates its alert settings. For example, if a new file extension triggers an alert and you mark it as a false positive, you will not receive an alert the next time the system sees that file extension.

Unsupported configurations

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

- NVMe volumes
- iSCSI volumes with ONTAP versions earlier than 9.17.1

Enable ARP/AI for a file system or a volume

Turning on ARP/AI for a file system automatically protects all current and new NAS volumes (NFS/SMB). Starting with ONTAP 9.17.1, it also supports SAN volumes (iSCSI). You can also turn on ARP/AI for specific volumes.

If ARP/AI detects a real attack after it is enabled, Workload Factory automatically creates a snapshot policy. This policy can create 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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From the **FSx for ONTAP** menu, select **Inventory** and then select the actions menu of the file system to enable ARP/AI and then select **Manage**.
4. 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.
5. 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.
6. Accept the statement to proceed.
7. Select **Apply** to save the changes.

Enable ARP/AI for a volume

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Volumes explorer**.
3. Select the actions menu of the volume to enable ARP/AI, then **Data protection actions**, and then **Manage ARP/AI**.
4. 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.
5. Accept the statement to proceed.
6. 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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Storage > FSx for ONTAP**.
3. From the **FSx for ONTAP** menu, select **Inventory** and then select the file system to validate ransomware attacks for.
4. From the file system overview, select the **Volumes** tab.
5. Select **Analyze attacks** from the Autonomous Ransomware Protection tile.
6. Download the attack events report to review if any files or folders were compromised and then decide if an attack has occurred.
7. If no attack occurred, select **False alarm** for the volume in the table and then select **Close**.
8. 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 the system detects a possible attack, it creates and locks a snapshot of the volume at that moment.

If the attack is later confirmed, you can restore affected files or the entire volume from the snapshot.

Locked snapshots cannot be deleted until the retention period expires. If the attack is later identified as a false alarm, the locked snapshot is deleted.

Because the system identifies the affected files and the time of the attack, you can recover only the affected files from available snapshots instead of restoring the entire volume.

Steps

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

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

3. From the **FSx for ONTAP** menu, select **Inventory** and then select the file system to recover data for.
4. From the file system overview, select the **Volumes** tab.
5. Select **Analyze attacks** from the Autonomous Ransomware Protection tile.
6. If an attack has occurred, select **Real attack** for the volume in the table.
7. 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.
8. 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 to the [NetApp Console](#).
2. Select the menu , **Workloads > Volumes explorer**.
3. Select the actions menu of the volume to clone.
4. Select **Data protection actions**, then **Clone volume**.
5. In the Clone volume dialog, enter a name for the volume clone.
6. Select **Clone**.

Protect your FSx for ONTAP data with a cyber vault

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

create immutable snapshots on the cyber vault volume.

What is a cyber vault?

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

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

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

FSx for ONTAP cyber vaults

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

Implementation

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

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

Related information

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

Manage FSx for ONTAP in Workload Factory

Analyze FSx for ONTAP 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) and latency events, leveraging Agentic AI with Amazon Bedrock integration.

About this task

Storage administrators often respond to FSx for ONTAP 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 the following event types from FSx for ONTAP file systems:

- **EMS events:** Error, alert, and emergency events from the ONTAP Event Management System.
- **Latency events:** Read and write latency events. Latency events can be viewed without a link connection.

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 and latency events, reducing manual investigation time.
- **Best-practice remediation:** The event analyzer gives clear, actionable steps to resolve FSx for ONTAP EMS and latency 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



If your account admin has turned off AI features, you'll see a message and won't be able to use AI analysis. Contact your administrator or [learn about AI administration](#).

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 view latency events you must do the following:

- [grant view, planning, and analysis permissions](#) in your AWS account.
- [Connect to the file system using a link](#).

To analyze EMS and latency events and receive AI-powered insights and recommendations, ensure the following:

- You have [operations and remediation permissions](#) in your AWS account.
- You have [connected to the file system using a link](#).
- 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 event analysis):

- **Amazon Bedrock activation**

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 events for FSx for ONTAP

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

Steps

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

Only FSx for ONTAP file systems that have events are shown.

5. Select the FSx for ONTAP file system you want to view.

The events table shows all events for that file system.

6. Review the event details in the table.
7. To analyze events, do one of the following:
 - To analyze multiple events, select the checkboxes next to the events, then select **Analyze** in the table.
 - To analyze one event and view insights in the side panel, select the event name, and then select

Analyze in the panel. The insights and recommendations are shown in the panel.

- To analyze one event from the table, select **Analyze** in the event row, then select the event name to view the insights and recommendations.

If needed, complete the AI analysis setup requirements by following the on-screen prompts to meet any missing prerequisites.

Related information

- [Set volume latency thresholds in Workload Factory](#)

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 supported cloud storage 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 configuration issues by category and a list of configuration issues and recommendations. You can review the recommendations for configuration issues. Workload Factory can automatically fix some issues; other issues 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.

You can also create custom rules in plain language to validate your environments against your organization's standards and NetApp best practices. Test rules with a dry run, then schedule them to run across your environments. [Learn more about custom rules.](#)

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 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

The well-architected dashboard assesses storage configurations against ONTAP best practices and 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.

- **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.

- **Snapshot policy**

Schedule local snapshots for efficient backup and quick restores. Snapshots are instant, point-in-time images of your volumes. This recommendation doesn't apply to database environments.

- **Application-consistent snapshots**

For database environments, use application-consistent snapshots with NetApp SnapCenter to take accurate, reliable snapshots of your volume data at a specific moment in time. This keeps your apps running smoothly and your data safe. SnapCenter makes backups easier and helps you restore data quickly and correctly, reducing downtime and protecting your most important workloads.

- **Cross-region replication**

Cross-region replication ensures that your data is replicated to another AWS region, providing enhanced data durability and availability. Set up cross-region replication to support disaster recovery and compliance.

- **SSD capacity headroom**

The SSD storage tier capacity should not exceed 80% utilization. This might impact data reads and writes to your capacity pool storage tier and impact the throughput capacity of your file system. When capacity runs out, data volumes become read-only. Services trying to write new data also fail.

- **Replication policy labels**

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

- **File object capacity**

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 maintaining file capacity utilization below 80% to allow new file creation.

Security

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

- **Ransomware protection (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 iSCSI volume access**

Volumes serving application data using iSCSI should not allow NAS access at the same time. Workload Factory recommends restricting iSCSI volumes from using any other protocol.

Operational excellence

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

- **Automatic capacity management**

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

- **Volume capacity headroom**

Keep volume utilization below 80% to prevent potential application read/write impacts. You can increase capacity manually or with volume autogrow.

- **FlexCache write mode**

For optimal performance, set the cache relationship write mode to best suit 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.

- **FlexCache replica size**

Enable volume autosize and scrubbing on cache volumes to maintain optimal size and focus the cache on hot data for peak efficiency.

- **Host-visible capacity reporting**

Enable host-visible capacity reporting to provide better visibility into storage usage at the volume level.

- **Operating System (OS) type**

Confirm that the ONTAP LUN operating system (OS) type value matches the operating system partitioning scheme to achieve I/O alignment. Incorrect configuration might reduce performance.

- **Block device space management**

Configure block device space settings for LUNs used by database host instances to prevent write failures and improve space efficiency on FSx for ONTAP. This configuration applies to the recommended combination of settings for thin-provisioned volumes: space reservation, space allocation, and fractional reserve.

- **Thin provisioning**

Configure thin provisioning for FSx for ONTAP volumes hosting database instances. This approach allows more logical data to be stored than physically available.

Cost optimization

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

- **Cold data tiering**

Enable cold data tiering to reduce SSD storage tier utilization. Apply a tiering policy to every volume focused on general data storage. For database volumes, enable cold data tiering where appropriate. FSx for ONTAP continuously scans the data for cold data and moves it to the capacity pool tier without disruption.

- **Storage efficiencies**

Enable storage efficiencies — compaction, compression, and deduplication — to optimize storage utilization and reduce the SSD tier cost.

- **Backup expiration**

Delete any manually created backups for volumes that were deleted, and any backups older than two weeks. These backups don't expire automatically, can increase costs, and add complexity if you no longer need them for SLA or recovery.

- **Snapshot expiration**

Delete old or manually created snapshots you no longer need. Removing them can lower costs and reduce clutter, especially if they aren't required for your SLA or recovery needs.

- **Inactive block devices**

After a block device isn't used for seven days, archive block device data or delete the unused block device to reduce costs.

- **Overprovisioned SSD capacity**

Excess SSD capacity adds unnecessary cost. Reduce the SSD tier while maintaining a 20% free space buffer. The operation may take hours to days, and no other file system operations are supported during that time.

- **Inactive NAS volumes**

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

Performance efficiency

Performance efficiency helps you ensure that your storage resources are used effectively and that workloads perform optimally.

- **FlexGroup volumes rebalance**

Rebalance your FlexGroup volumes to improve efficiency. After scaling out, capacity and workload activity can become uneven across the volume.

- **FlexVol volumes rebalance**

Rebalance your FlexVol volumes to improve efficiency. After scaling out, capacity and workload activity can become uneven across the volume.

What's next

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

Implement well-architected FSx for ONTAP file system configurations

Use 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, and fix issues to improve the architecture of any systems that aren't optimized for reliability, security, efficiency, performance, and cost.

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

About this task

Workload Factory analyzes Amazon FSx for NetApp ONTAP file system deployment configurations daily.

Because requirements for storage configurations vary, you can dismiss the analysis of specific configurations that don't apply to your storage environment.

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.

What is analyzed

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

- Reliability: volume backups, snapshot policy, cross-region replication, SSD capacity headroom, replication policy labels, and file object capacity
- Security: Ransomware Protection (ARP/AI) and unauthorized iSCSI volume access
- Operational excellence: automatic capacity management, volume file capacity headroom, FlexCache write mode, FlexCache replica size, host-visible capacity reporting, OS type, block device space management, and thin provisioning
- Cost optimization: cold data tiering, storage efficiencies, backup expiration, snapshot expiration, inactive block devices, overprovisioned SSD capacity, and inactive NAS volumes
- Performance efficiency: FlexGroup volumes rebalance and FlexVol volumes rebalance

Before you begin

- You must [grant operations and remediation permissions](#) in your AWS account.
- [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.

Fix operations can cause instance downtimes or service interruptions. Review each recommendation carefully before fixing a configuration issue.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads** > **Well-architected** > **FSx for ONTAP**.
3. Select **View and fix** for any configuration. Make sure you review the recommendation carefully.

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

4. Select resources and then **Fix**.

Some fix operations are only available for specific operations and not always for more than one resource. For example, you might be able to fix the configuration for a single volume but not for multiple volumes.

5. Review the action items to learn about consequences of this fix such as instance downtimes or service interruptions.
6. Select **Fix** if you want Workload Factory to apply the recommended changes.

Result

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

Dismiss a configuration analysis

Stop a configuration analysis indefinitely for one or more resources. You can restart the analysis when needed.

Dismiss to stop configuration analysis for one or more FSx for ONTAP file systems or volumes. You can restart the analysis anytime.

Dismiss a configuration analysis for configurations

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads** > **Well-architected** > **FSx for ONTAP**.
3. Under Configurations, identify the configuration that doesn't apply to your environment.
4. Select **...** and then select **Dismiss**.
5. In the Dismiss configuration dialog, select **Dismiss** to stop the analysis for the configuration.

Dismiss a configuration analysis for your resources

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Well-architected > FSx for ONTAP**.
3. In the FSx for ONTAP well-architected dashboard, under Configurations, select the configuration that doesn't apply to your environment.
4. On the configuration page, select the resource(s) to dismiss from the configuration analysis.
 - For one resource: select  and then select **Dismiss**.
 - For multiple resources: select the resources and then select **Dismiss** next to Bulk action.
5. In the Dismiss resources dialog, select **Dismiss** to stop the analysis.

Result

Workload Factory stops the configuration analysis for selected configurations.

Reactivate a dismissed configuration analysis

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

Reactivate a configuration analysis for configurations

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Well-architected > FSx for ONTAP**.
3. Select **Dismissed configurations**.
4. Identify the configuration(s) you want to reactivate and select **Reactivate**.

Reactivate a configuration analysis for your resources

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Well-architected > FSx for ONTAP**.
3. In the FSx for ONTAP well-architected dashboard, under Configurations, select the configuration that doesn't apply to your environment.
4. On the configuration page, select **Dismissed configurations**.
5. Select the resource(s) to reactivate.
 - For one resource: select  and then select **Reactivate**.
 - For multiple resources: select the resources and then select **Reactivate** next to Bulk action.

Result

Workload Factory reactivates the configuration analysis and runs a new analysis every day.

Custom rules

Learn about custom rules in NetApp Console

Custom rules let you set, check, and enforce your own architecture standards, along with NetApp best practices, in the Well-architected dashboard.

Overview

The Rule builder lets you create custom best-practice rules using plain language. These rules are checked across all of your resources.

Custom rules help you enforce your organization's standards, risk limits, and operational requirements—things that generic vendor rules might not fully cover.



If your account administrator has turned off AI features, the Rule builder will not be available and you will not be able to create custom rules. [Learn about managing AI features.](#)

How custom rules work

Custom rules use an AI-guided process that turns your plain-language instructions into clear, structured rules. The AI acts as a translator and checker. It suggests a rule, explains how it understood your request, and asks for your approval before the rule is saved or scheduled.

This process ensures that:

- The rule logic is clearly defined and checked for errors
- No free-form code is run
- You can see exactly what will be evaluated before you approve the rule

Limitations

The following limitations apply to custom rules:

- Automatic fixes are not available. You must correct non-compliant resources yourself using the guidance provided.
- Each rule can check only one type of resource. To check multiple resource types, create separate rules.
- The shortest evaluation interval is 1 hour to avoid excessive scanning.
- Rules have read-only access and cannot make changes to your environment.

Rule definition model

Each custom rule consists of the following components:

Rule name

A user-defined name for the rule.

Description

Your rationale for why this rule matters to your organization.

Resource type

The FSx resource to evaluate (file system, volume, or cache). Currently, one resource type per rule.

Scheduler

How often the rule is evaluated (minimum 1 hour, default 24 hours).

Condition logic

The specific criteria to check.

Evaluation scope

Which resources are evaluated (all applicable resources in scope).

Remediation logic

Manual guidance for addressing non-compliant resources.

Rule lifecycle

Custom rules progress through the following states:

Draft

Rule is being created. No evaluation occurs.

Dry-run

One-time evaluation to preview which resources would be affected. Used for validation before scheduling.

Scheduled (Enabled)

Rule is active and runs on the defined schedule.

Disabled

Rule is paused but not deleted. Can be re-enabled.

Deleted

Rule is permanently removed.

Custom rule examples

The following examples illustrate common custom rules:

Production volume protection

"All volumes with tag 'Prod' must have a snapshot policy installed"

This rule automatically creates snapshots to protect production volumes according to company requirements.

Archive volume tiering

"All volumes with tag 'Archive' should have a tiering policy set to 'All'"

This rule lowers costs by automatically moving archived data to cheaper storage.

File system naming standards

"All file systems must follow the naming convention 'fsx-{environment}-{application}'"

This rule validates that file systems adhere to organizational naming conventions for easier management and cost allocation.

What to do next

[Create a custom rule](#)

Create a custom rule in NetApp Workload Factory

Create custom rules by describing them in plain language. The AI rule builder in NetApp Workload Factory turns your request into a validated rule that checks your FSx for ONTAP resources.

About this task

The rule creation process walks you through the steps: defining the rule, reviewing how the AI interpreted it, and checking the results before scheduling regular evaluations. You stay in control at every step. The system suggests a rule based on your prompt, but nothing moves forward until you confirm it.



If your account administrator turned off AI features, you won't be able to create custom rules. Contact your administrator or [learn about AI administration](#).

Limitations

The following limitations apply to custom rules:

- Automatic fixes are not available. You must correct non-compliant resources yourself using the guidance provided.
- Each rule can check only one type of resource. To check multiple resource types, create separate rules.
- The shortest evaluation interval is 1 hour to avoid excessive scanning.
- Rules have read-only access and cannot make changes to your environment.

Custom rule examples

The following examples illustrate common custom rules:

Production volume protection

"All volumes with tag 'Prod' must have a snapshot policy installed"

This rule automatically creates snapshots to protect production volumes according to company requirements.

Archive volume tiering

"All volumes with tag 'Archive' should have a tiering policy set to 'All'"

This rule lowers costs by automatically moving archived data to cheaper storage.

File system naming standards

"All file systems must follow the naming convention 'fsx-{environment}-{application}'"

This rule validates that file systems adhere to organizational naming conventions for easier management and cost allocation.

Before you begin

Before you can create and run custom rules, ensure you have the following:

- You have configured AWS credentials with read-only access. [Add AWS credentials to Workload Factory](#).
- You have ONTAP credentials with viewer permissions. [Add ONTAP credentials to Workload Factory](#).
- You have a link or Console agent configured connectivity to the cloud storage system. [Create and associate a link](#) or [deploy a Console agent](#).
- You understand that custom rule creation uses NetApp's Bedrock AI service.

Create a rule

Use the Rule builder to create custom best-practice rules using plain language. These rules are checked across all of your resources.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Well-architected**.
3. From the Well-architected menu, select **Rule builder**.

When you open the rule builder, it shows sample prompts to help you write your custom rule. You can select **Copy prompt** to copy a sample prompt into the AI dialog text box.

4. In the **Create a rule** field, enter your rule description in natural language. Be as specific as possible about:
 - What resource type you're checking (file system, volume, or cache).
 - What condition you're validating.
 - What threshold defines compliance.
 - Why this rule matters (optional but recommended).
5. Select .

The AI analyzes your prompt and generates a rule definition.

Review the AI interpretation

The rule builder displays:

- **Parsed interpretation:** How the AI understood your prompt in plain language.
- **Rule definition:** The structured rule including resource type, condition logic, and evaluation criteria.
- **Reasoning:** Why the AI interpreted your prompt this way.
- **Confidence indicator:** How confident the system is in the interpretation.

Steps

1. Review the interpretation to make sure it matches what you intended.

Check that the rule definition is correct.

If the interpretation is unclear or wrong, the system may ask questions such as:

- * Which resource should be checked?
- * What value or limit defines compliance?

* How often should the rule run?

2. Respond to the AI's clarifying questions.

The AI might ask you to choose from several options to better understand what you need. Answer each question and select **Continue**, or select **Skip** to move past optional questions.

3. After answering all questions, select **Continue** to create your rule.

Test the rule

Before scheduling the rule, run a one-time test to see which resources it will affect.

Steps

1. Select **Validate rule**.

The system evaluates all resources in scope against your rule logic.

2. Review the validation results:
 - **Resources evaluated:** Total number of resources checked
 - **Compliant:** Resources that meet the rule criteria
 - **Non-compliant:** Resources that do not meet the criteria
 - **Failed evaluations:** Resources that could not be evaluated (with error details)
3. If some resources failed evaluation, review the error reasons:
 - **AWS credential errors:** Read-only permissions may be insufficient for certain resources
 - **Connectivity errors:** Link or agent might not be configured for specific file systems
 - **ONTAP permission errors:** The viewer user might not have access to certain resources

Save or schedule the rule

After you test your rule, you can save it as a draft or set it to run regularly.

Steps

1. Do one of the following:
 - To schedule regular evaluation, set the evaluation frequency (minimum 1 hour, default 24 hours) and select **Schedule**. The rule is enabled and runs on the defined schedule.
 - To save without scheduling, select **Save as draft**. The rule is saved but not evaluated until you schedule it.
 - To discard the rule, select **Discard** to cancel rule creation and start over.

Result

Your custom rule has been created and is listed in the rule catalog. If you set a schedule, it will check your FSx for ONTAP resources at the specified frequency.

Troubleshooting

No file systems available

Set up your AWS credentials and make sure your environment has at least one FSx for ONTAP file system. The rule builder needs a file system to apply rules.

Low confidence interpretation

More information is needed. Answer the questions or rewrite your prompt with clear details about the resource type, limits, and conditions.

Partial validation failures

Check which resources failed and why. Common causes are missing link or agent settings, or not having enough permissions. Fix the problem, then validate again.

Ambiguous prompt

Be more specific about what you want to check. Include the resource type, the condition, and the threshold. See the prompt examples for help.

Rule targets multiple resource types

Create a separate rule for each resource type. Only one resource type per rule is supported.

What to do next

[Manage custom rules](#)

Manage custom rules for NetApp Workload Factory

The Rules catalog is one place to manage all your custom rules. You can see each rule's details, check its results, and turn it on or off.

View the Rules catalog

The Rules catalog shows all the rules you created, their status, schedule, and latest results.

Steps

1. Log in using one of the [console experiences](#).
2. Select the menu  and then select **Rule builder**.
3. Click **View your rules** to open the Rules catalog.

View rule results

The results show which resources follow your rule and which do not.

Steps

1. In the Rules catalog, click the rule you want to review.
2. Select the **All runs results** tab.
3. Select **View details** to see the results for each run.

If the check failed for some items, read the error message to see why. Common reasons are connection problems, missing permissions, or a link that hasn't been set up.

Edit a custom rule

You can change a rule's logic, schedule, or other settings at any time. When you edit a rule, the rule builder opens with the rule's current settings.

Steps

1. In the Rules catalog, click the rule you want to edit.
2. Click **Edit**.

The rule switches to draft mode and the rule builder opens.

3. Edit the rule prompt or use the conversation to improve the rule.
4. Review the AI's updated interpretation and rule definition.

The system shows what changed from the current rule, including any API calls or logic.

5. Do one of the following depending on the changes made:
 - If minor changes were made, click **Save** to update the rule.
 - If significant changes were made, click **Validate rule** to test the changes, then click **Schedule** to save and schedule the rule, or click **Save as draft** to save without scheduling.
 - To discard the changes, click **Revert to previous version**. The rule returns to its original definition and leaves edit mode.

5

Deploy the recommended Google Cloud NetApp Volumes NFS datastores

Deploy the new Google Cloud NetApp Volumes storage pools and volumes that will provide the external NFS datastores for your VMs in your Google Cloud VMware Engine environment.

[Learn how to deploy your Google Cloud NetApp Volumes NFS datastores.](#)

What's next

After you deploy the Google Cloud NetApp Volumes datastores, you can migrate your VMs from your on-premises environment to your Google Cloud VMware Engine cluster and begin taking advantage of the cost savings and flexibility provided by decoupling storage from compute.

What's next

You can [deploy the Google Cloud NetApp Volumes NFS datastores](#) using the plan details and Terraform automation code.

Depending on the policy and permissions that you added to your Workload Factory account, you can deploy the FSx for ONTAP file system completely using Workload Factory (using Automate mode). If you have fewer permissions (Read mode), or no permissions (Basic mode), you'll need to use the CloudFormation information from the Codebox and deploy the FSx for ONTAP file system yourself in AWS.

Requirements for deployments to Amazon EC2

- The root volume of the FSx for ONTAP file system must be created as an EBS volume.

Steps

1. At the bottom of the **Review plan** page, select **Deploy** and the Create an FSx for ONTAP file system page is displayed.

Most of the fields that define your FSx for ONTAP file system are completed based on the information you provided, but there are a few fields that you need to complete in this page.

You can use the Quick create or Advanced create option. Advanced create offers a few additional storage parameters that you can customize. [Learn what these two options offer.](#)

2. **AWS credentials:** Select or add credentials that will give Workload Factory the permissions necessary to create your FSx for ONTAP file system directly. You can also select the CloudFormation code from Codebox and deploy the FSx for ONTAP file system yourself in AWS.
3. **File system name:** Enter the name that you want to use for this FSx for ONTAP file system.
4. **Tags:** Optionally you can add tags to categorize this FSx for ONTAP file system.
5. In the "Network & security" section, enter the following information:

- a. **Region & VPC:** Select the Region and the VPC where the FSx for ONTAP file system will be deployed.

If you are deploying to VMware Cloud on AWS, ensure you deploy it in a VPC that is different than the VPC where the VMware Cloud on AWS is deployed.

- b. **Security group:** When using the **Advanced create** option, you can select the default security group for the FSx for ONTAP VPC so that all traffic can access the FSx for ONTAP file system.

You can add an inbound rule that restricts what other AWS services can access the FSx for ONTAP file system. This will block the amount of services that are open. These are the minimum ports and protocols:

Protocols	Ports	Purpose
TCP, UDP	111	Portmapper (used to negotiate which ports are used in NFS requests)
TCP, UDP	635	NFS mountd (receives NFS mount requests)
TCP, UDP	2049	NFS network traffic
TCP, UDP	4045	Network Lock Manager (NLM, lockd) - Handles lock requests.
TCP, UDP	4046	Network Status Monitor (NSM, statd) - Notifies NFS clients about reboots of the server for lock management.

- c. **Availability zone:** Select the Availability Zone and the Subnet.

You should select the same availability zone as where your VMware SDDC is deployed if you want to avoid charges for cross-AZ traffic.

- d. **Encryption:** When using the **Advanced create** option, you can select the AWS encryption key name from the dropdown.
- e. **Datastore access control:** When using the **Advanced create** option, you can select whether all hosts can access the datastores or whether only certain vSphere cluster nodes on a specific subnet can access the datastores.

6. In the "File system details" section, enter the following information:

- a. **ONTAP credentials:** Enter and confirm the ONTAP password.
- b. **Storage VM credentials** (Advanced create only): Enter and confirm the storage VM password. The password can be specific to this file system, or you can use the same password entered for ONTAP credentials.

7. In the **Summary** section, you can view the FSx for ONTAP file system and datastore configuration that the VMware migration advisor has designed based on your information.

8. Select **Create** to deploy the FSx for ONTAP file system. This process can take up to 2 hours.

Optionally, in the Codebox window you can select **Redirect to CloudFormation** to create the file system using a CloudFormation stack.

In either case, you can monitor the creation progress in CloudFormation.

Result

The FSx for ONTAP file system is deployed.