



Use Amazon FSx for NetApp ONTAP

Amazon FSx for NetApp ONTAP

NetApp

February 02, 2026

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

Table of Contents

Use Amazon FSx for NetApp ONTAP	1
Explore savings with FSx for ONTAP in NetApp Workload Factory	1
Calculator options	1
Explore savings via customization	1
Explore savings for detected storage environments	3
Deploy FSx for ONTAP file systems	5
Track costs for your resources in NetApp Workload Factory	5
Use links	6
Learn about NetApp Workload Factory links	6
Connect to an FSx for ONTAP file system with a Lambda link	7
Manage Workload Factory links	14
Discover cache volumes in Workload Factory	17
Manage volumes	18
Create an FSx for ONTAP volume in Workload Factory	18
Access your FSx for ONTAP file system data	23
Create block storage resources	24
Create an initiator group for a file system in NetApp Workload Factory	24
Create a block device for a file system in NetApp Workload Factory	25
Create a storage VM for an FSx for ONTAP file system	27
Create a storage VM	27
Protect your data	28
Types of data protection in NetApp Workload Factory	28
Use snapshots	30
Use backups to object storage	33
Use replication	35
Protect your data with NetApp Autonomous Ransomware Protection with AI	38
Clone a volume in NetApp Workload Factory	41
Use on-premises ONTAP cluster data in NetApp Workload Factory	42
Protect your data with a cyber vault	45

Use Amazon FSx for NetApp ONTAP

Explore savings with FSx for ONTAP in NetApp Workload Factory

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

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

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

Calculator options

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

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

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

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

Explore savings via customization

Follow the steps under the tab for your storage selection.

Amazon Elastic Block Store (EBS)

Steps

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

Amazon FSx for Windows File Server

Steps

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

Amazon Elastic File System (EFS)

Steps

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

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

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

- [Export PDF](#)
- [Send by email](#)
- [View the calculations](#)

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

Explore savings for detected storage environments

Before you begin

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



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

Follow the steps under the tab for your storage selection.

Amazon Elastic Block Store (EBS)

Steps

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

Amazon FSx for Windows File Server

Steps

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

Amazon Elastic File System (EFS)

Steps

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

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

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

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

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

Deploy FSx for ONTAP file systems

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

Deployment methods

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

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

Track costs for your resources in NetApp Workload Factory

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

About this task

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

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

Before you begin

Grant credentials with the *view, planning, and analysis* permission policy in Workload Factory to track FSx for ONTAP costs.

Steps

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

Use links

Learn about NetApp Workload Factory links

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

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

Why links are beneficial

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

Here are some benefits of using links:

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

How links work

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

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

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

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

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

Costs

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

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

When a link is required

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

The features that require a link include:

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

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

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

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

[Learn more about links.](#)

About this task

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

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

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



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

Associate a new link

Associating a new link includes link creation and association.

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

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

Before you begin

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

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

Create automatically

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

Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the actions menu of the file system to associate a link to and then select **Associate link**.
5. In the Associate link dialog, select **Create a new link** and select **Continue**.
6. On the Create Link page, provide the following:
 - a. **Link name**: Enter the name that you want to use for this link. The name must be unique within your account.
 - b. **AWS Secrets Manager**: Optional. Allows Workload Factory to fetch FSx for ONTAP access credentials from your AWS Secrets Manager.

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

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

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

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

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

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

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

Enter the Lambda execution role ARN in the text box.

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

by FSx for ONTAP file systems.

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

7. Select **Create**.

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

8. Select **Continue** to open the AWS Management Console, and then log in to the AWS account for this FSx for ONTAP file system.
9. On the Quick create stack page, under Capabilities, select **I acknowledge that AWS CloudFormation might create IAM resources**.

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

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

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

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

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

Create manually

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

Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the actionsenu of the file system to associate a link to and then select **Associate link**.
5. In the Associate link dialog, select **Create a new link** and select **Continue**.
6. On the Create Link page, select CloudFormation or Terraform from the Codebox, and then provide the following:
 - a. **Link name**: Enter the name that you want to use for this link. The name must be unique within your account.
 - b. **AWS Secrets Manager**: Optional. Allows Workload Factory to fetch FSx for ONTAP access credentials from your AWS Secrets Manager.

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

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

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

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

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

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

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

Enter the Lambda execution role ARN in the text box.

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

e. **Link registration:** Select CloudFormation or Terraform for the instructions for how to register the link, and follow the instructions.

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

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

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

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

7. Select **Create**.

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

minutes.

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

Result

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

Associate an existing link with an FSx for ONTAP file system

After you create a link, associate it with one or more FSx for ONTAP file system.

Steps

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

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

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



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

7. Select **Apply**.

Result

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

Troubleshoot issues with AWS Secrets Manager link authentication

Issue

The link lacks permissions to retrieve the secret.

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

Issue

The secret isn't found.

Resolution: Provide the correct secret ARN.

Issue

The secret isn't in the right format.

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

The secret should contain the following key valid pairs:

- filesystemID = FSx_filesystem_id
- username = FSx_user
- password = user_password

Issue

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

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

Manage Workload Factory links

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

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

View the links associated with your account


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

Steps

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

Existing links appear on the Links page.

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



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

Edit a link

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



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

Authenticate a link

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

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



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

Steps

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

3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the actions menu of the file system to associate a link to and then select **Manage**.
5. In the file system overview, select **Authenticate the link**.
6. In the Authenticate link page, select an authenticate mode:
 - Workload Factory: enter the password twice.
 - AWS Secrets Manager: enter the secret ARN.
7. Select **Apply**.

Result

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

Update the password for link authentication

When the administrative password is invalid, update the password to connect the link to the FSx for ONTAP file system.

Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the actions menu of the file system to associate a link to and then select **Manage**.
5. In the file system overview, select **Update password**.
6. In the Authenticate link page, enter the new password twice.
7. Select **Apply**.

Result

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

Remove a link

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

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

Steps

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

Existing links appear on the Links page.

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

Links		
Link (1)		
 Link name	Lambda Type us-east-1 vpc-123456 10.2.3.0/24 Location	1 View Associated resources arn:4447575775767676868686.. ARN
Connected Status account-1234666495 AWS account	Lambda Type us-east-1 vpc-123456 10.2.3.0/24 Location	

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

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

Discover cache volumes in Workload Factory

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

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

About this task

Link association is required to discover cache relationships.

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

Before you begin

Consider the following before you begin.

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

Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the **Cache relationships** tab.

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

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

Related information

[Manage cache volumes](#)

Manage volumes

Create an FSx for ONTAP volume in Workload Factory

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

About this task

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

The following protocols might be used for your volumes:

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

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

Details for volume settings

Immutable files

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

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

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



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

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

Before you begin

Review the following prerequisites before you create a volume:

- You must have an FSx for ONTAP file system in the Workload Factory console.
- You must have a storage VM.
- For protocol access, complete the following:
 - To configure access to the volume, you must associate a link. [Learn how to associate an existing link or to create and associate a new link](#). After the link associates, return to this operation.
 - You must configure access for the protocol you select, either SMB/CIFS, NFS, or iSCSI.

Create a volume

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



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

Steps

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

2. Select the menu and then select **Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the actions menu of the file system you want to create a volume in, and select **Create volume**.
5. On the Create volume page under General details, provide the following details:
 - a. **Volume name**: Enter a name for the volume.
 - b. **Storage VM name**: Optionally, enter a storage VM name.
 - c. **Volume style**: Select **FlexVol** or **FlexGroup** volume.

FlexVol volume style is selected by default.

FlexGroup volume size depends on the number of constituents, requiring 100 GiB per constituent.
 - 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.
 - e. **Volume autogrow**: Optionally, enable volume autogrow to automatically expand volume capacity until the volume reaches the maximum size. This feature accommodates increasing data usage, ensuring uninterrupted operations.

Specify the maximum volume growth size and unit. You cannot set the maximum growth size smaller than the current volume size
 - f. **Tags**: Optionally, you can add up to 50 tags.
6. Under Access (only for file systems with associated links), provide the following details:
 - a. **Access type**: Select **File access** or **Block access**. Additional fields to configure volume access differ depending on your selection.
 - **File access**: allows multiple authorized users and devices access to the volume using SMB/CIFS, NFS, or dual (SMB/NFS) protocols.

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

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

If you selected an existing export policy, enter a new policy name.
 - iii. **Add Export Policy Rule**: Optionally, for a custom export policy, you can add export policy rules to the policy.
 - c. **SMB/CIFS share**: Provide the following:
 - i. **Name**: Enter the SMB/CIFS share name to provide access.
 - ii. **Permissions**: Select Full control, Read/Write, Read, or No access, and then enter the users or groups separated by a semicolon (;). Users or groups are case sensitive and the user's domain

must be included using the format "domain\username".

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

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

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

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

- Select **Create a new initiator group** or **Map an existing initiator group**.
- Select the **Host operating system** from the dropdown menu.
- Enter an **Initiator group name** for a new initiator group.
- 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 metadata:** Enabling metadata creates an S3 table containing all objects accessible by the S3 access point, which you can use for auditing, governance, automatic, analysis, and optimization. Enabling metadata incurs additional AWS costs. Refer to [Amazon S3 pricing documentation](#) for more information.

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

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

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

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

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

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

default

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

- A maximum of six hourly snapshots taken five minutes past the hour.

- A maximum of two daily snapshots taken Monday through Saturday at 10 minutes after midnight.
- A maximum of two weekly snapshots taken every Sunday at 15 minutes after midnight.



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

default-1weekly

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

none

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

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

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

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

- 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 iSCSI and NVMe protocols are not supported for ARP/AI.
- The file system already has an ARP/AI policy.

8. Under Advance configuration, provide the following:

- Junction path:** Enter the location in the storage VM's namespace where the volume gets mounted. The default junction path is /<volume-name>.
- 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.

9. Select **Create**.

Related information

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

Access your FSx for ONTAP file system data

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

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

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

Get mount point for volumes in NetApp Workload Factory

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

Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From **FSx for ONTAP**, select the actions menu of the file system with the volume and then select **Manage**.
4. From the file system overview, select the **Volumes** tab.
5. From the Volumes tab, select the actions menu for the volume, then **Basic actions**, and then **View mount command**.
6. In the Mount command dialog, select **Copy** to copy the command for either the NFS mount point or CIFS share. You'll enter the copied command in your terminal.

7. Select **Close**.

Connect to NAS clients

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

Connect to SAN clients

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

Create block storage resources

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

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

About this task

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

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

Before you begin

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

Steps

1. Log in using one of the [console experiences](#).
2. In the Storage tile, select **Go to Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the **Block devices** tab.
5. Select the resource type **Create initiator group** and then select **Create igroup**.
6. In the **Create initiator group** dialog, do the following:
 - **igroup name**: Enter a name for the initiator group.
 - **igroup description**: (Optional) Enter a description for the initiator group.
 - **Storage VM name**: Select the storage VM for the initiator group.
 - **Block device name**: Select one or more block devices to associate with the initiator group. The block devices listed are those that have not been mapped to a host initiator yet.

- **Operating system type:** Select Linux, VMware, or Windows for the operating system type.
- **Host initiators:** Add one or more iSCSI qualified (IQN) host initiators to the initiator group.

7. Select **Create**.

Related information

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

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

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

About this task

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

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

Before you begin

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

Steps

1. Log in using one of the [console experiences](#).
2. In the Storage tile, select **Go to Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the **Block devices** tab.
5. Select **Create block device**.
6. Under **Volume details**, do the following:
 - a. **Volume name:** Select one of the following options.
 - Create a new volume and enter the name of the volume.
 - Select an existing volume.
 - b. **Storage VM:** Select a storage VM.
 - c. **Volume style:** The default volume style is **FlexVol**.
 - d. **Volume size:** Enter the size of the volume and select the unit. The maximum size per FlexVol volume is 100 TiB.
 - e. **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.
 - f. **Tags:** Optionally, add tags to help organize and categorize your block device.
7. Under **Block device details**, do the following:
 - a. **Block device name:** Enter a name for the block device.
 - b. **Block device size:** Enter the size of the block device and select the unit. The block device size must be smaller than the available volume size.
8. Under **Access**, do the following:
 - a. **iSCSI configuration:** Select one of the following options.

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

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

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

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

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

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

default

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

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



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

default-1weekly

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

none

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

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

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

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

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

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

10. Select **Create**.

Related information

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

Create a storage VM for an FSx for ONTAP file system

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

About this task

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

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

Before you begin

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

Create a storage VM

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



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

Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the actions menu of the file system and then select **Manage**.
5. In the file system overview under Storage VMs, select **Create storage VM**.
6. On the Create storage VM page, under Storage VM configuration, provide the following:
 - a. **Name**: Enter a name for the storage VM.
 - b. **Storage VM credentials**: Provide a password for this storage VM's `vsadmin` user or use the file system's `fsxadmin` user credentials.
 - c. **Root volume security style**: Select the root volume security style depending on the type of clients that

access your data - UNIX (Linux clients), NTFS (Windows clients), or Mixed.

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

7. Select **Create**.

Protect your data

Types of data protection in NetApp Workload Factory

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

Types of data protection

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

Snapshots

A snapshot creates a read-only, point-in-time image of a volume within the source volume as a snapshot copy. You can use the snapshot copy to recover individual files, or to restore the entire contents of a volume. Snapshots are the basis of all the backup methods. The snapshot copy that is created on your volume is used to keep the replicated volume and backup file synchronized with changes made to the source volume.

NetApp Autonomous Ransomware Protection with AI

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

Replication


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

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

Backups

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

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



Best practices for protecting your workload data

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

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

Protect your workload data with snapshots

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

Protect your workload data with NetApp Autonomous Ransomware Protection with AI

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

Protect your workload data with volume replication

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

Protect your workload data with backups

Volume backups provide independent point-in-time copies of your volume. They can be used to store old backups and provide the necessary second copy of your data. Daily, weekly, and monthly backup schedules

allow for RPOs starting at one day. Volume backups can only be restored as a whole. Creating a volume from a backup (RTO) can take hours to many days, depending on the size of the backup.

Recommendations for protecting your workload data

Consider the following recommendations for protecting your workload data.

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

Use snapshots

Create a manual snapshot of an FSx for ONTAP volume

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

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

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



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

Before you begin

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

Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the actions menu of the file system that contains the volume to create a snapshot for and then select **Manage**.
5. In the file system overview, select the **Volumes** tab.
6. From the **Volumes** tab, select the actions menu for the volume to protect with snapshots.
7. Select **Data protection actions** and then **Manage snapshots**.

8. From the Manage snapshots page, select **Create snapshot**.
9. In the Create a snapshot dialog, do the following:
 - a. Enter a snapshot name in the **Snapshot name** field.
 - b. Optionally, select a label or create a new label.
 - c. Set the **Retention period** as a number of hours, days, months, or years.
 - d. Optional: **Make this snapshot immutable** to prevent the snapshot from being deleted during the retention period.

Accept the statement regarding immutable snapshots.

10. Select **Create**.

Create a snapshot policy for storage VMs in Workload Factory

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

About this task

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

- default
- default-1weekly
- none

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

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

Before you begin

- Once a snapshot policy is created, its association with the storage VM(s) cannot be modified, but you can always add or remove the policy from volumes.
- Consider the following about snapshot capacity before you use snapshots:
 - For most datasets, an additional capacity of 20% is enough to keep snapshots for up to four weeks. As data gets older, its use for restorations becomes less likely.
 - Overwriting all the data in a snapshot consumes significant volume capacity, which factors into provisioning volume capacity.
- To create a custom snapshot policy, you must associate a link. [Learn how to associate an existing link or to create and associate a new link](#). After the link associates, return to this operation.

Steps

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

3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the actions menu of the file system with the volume and then select **Manage**.
5. In the file system overview, select the **Storage VMs** tab.
6. From the **Storage VMs** tab, select the actions menu for the volume to protect with scheduled snapshots, then **Advanced actions**, and then **Manage snapshot policies**.
7. On the Snapshot policy management page, select **Create snapshot policy**.
8. In the **Snapshot policy name** field, enter a name for the snapshot policy.
9. Optionally, enter a description for the snapshot policy.
10. Under **Policy schedule and copies**, select when to create snapshots. For example, every minute or hourly.

You can select more than one frequency.

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

The maximum number of copies is 1,023.

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

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

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

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

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

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

16. Select **Create**.

Restore a volume from a snapshot in Workload Factory

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

About this task

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

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

Before you begin

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

- You can only restore a volume from a snapshot if you have an existing snapshot copy of the volume.

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

Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the actions menu of the file system with the volume and then select **Manage**.
5. In the file system overview, select the **Volumes** tab.
6. From the **Volumes** tab, select the actions menu for the volume to restore from a snapshot.
7. Select **Data protection actions** and then **Manage snapshots**.
8. From the Manage snapshots page, select the actions menu for the snapshot to restore, and then select **Restore**.
9. In the Restore volume from a snapshot dialog, select from the following options:
 - Toggle to select **Restore as a new volume**.

In the **Restored volume name** field, enter a unique name for the volume to restore.

 - Restore data from a snapshot to an existing volume. This operation permanently deletes any data that was modified after the snapshot creation time.

Accept the statement to proceed.
10. Select **Restore**.

Use backups to object storage

Create a manual backup of a volume in NetApp Workload Factory

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

About this task

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

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

Before you begin

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

Steps

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

2. Select the menu and then select **Storage**.
3. From the Storage menu, select **FSx for ONTAP**.
4. From **FSx for ONTAP**, select the actions menu of the file system with the volume and then select **Manage**.
5. In the file system overview, select the **Volumes** tab.
6. From the **Volumes** tab, select the actions menu for the volume to back up.
7. Select **Data protection actions**, **FSx for ONTAP backup**, and then **Manual backup**.
8. In the Manual backup dialog, enter a name for the backup.
9. Select **Back up**.

Restore a volume from a backup in NetApp Workload Factory

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

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

Steps

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

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

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

10. Select **Restore**.

Use replication

Create a replication relationship in NetApp Workload Factory

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

About this task

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

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

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

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

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

Before you begin

Consider the following before you begin.

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

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

Steps

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

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

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

a. **Use case**: Select one of the following use cases for the replication. Depending on the selected use case, Workload Factory fills in the form with recommended values in accordance with best practices. You can accept the recommended values or make changes as you complete the form.

- Migration: transfers your data to the target FSx for ONTAP file system
- Hot disaster recovery: ensures high availability and rapid disaster recovery for critical workloads
- Cold or archive disaster recovery:
 - Cold disaster recovery: uses longer recovery time objectives (RTO) and recovery point objects (RPO) to lower costs
 - Archive: replicates data for long-term storage and compliance
- Other

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

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

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

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

c. **Storage VM name**: Select the storage VM from the dropdown menu. The storage VM you select is the target for all selected volumes in this replication relationship.

d. **Volume name**: The target volume name is generated automatically with the following format {OriginalVolumeName}_copy. You can use the auto-generated volume name or enter another volume name.

e. **Tiering policy**: Select the tiering policy for the data stored in the target volume. The tiering policy defaults to the recommended tiering policy for the use case you selected.

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

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

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

Without a limit, network and application performance may decline. Alternatively, we recommend an

unlimited transfer rate for FSx for ONTAP file systems for critical workloads, for example, those that are used primarily for disaster recovery.

7. Under Replication settings, provide the following:

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

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

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



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

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

- c. **Immutable snapshots:** Optional. Select **Enable immutable snapshots** to prevent snapshots taken in this policy from being deleted during the retention period.
 - Set the **Retention period** in number of hours, days, months, or years.
 - **Snapshot policies:** In the table, select the snapshot policy frequency and the number of copies to retain. You can select more than one snapshot policy.
- d. **S3 access point:** Optionally, attach an S3 access point to access FSx for ONTAP file system data residing on NFS or SMB/CIFS volumes via AWS S3 APIs. Only the file access type is supported. Providing the following details:
 - **S3 access point name:** Enter the name of the S3 access point.
 - **User:** Select an existing user with access to the volume or create a new user.
 - **User type:** Select **UNIX** or **Windows** as the user type.
 - **Network configuration:** Select **Internet** or **Virtual private cloud (VPC)**. The type of network you choose determines whether the access point is accessible from the internet or restricted to a specific VPC.
 - **Enable metadata:** Enabling metadata creates an S3 table containing all objects accessible by the S3 access point, which you can use for auditing, governance, automatic, analysis, and optimization. Enabling metadata incurs additional AWS costs. Refer to [Amazon S3 pricing documentation](#) for more information.
- e. **S3 access point tags:** Optionally, you can add up to 50 tags.

8. Select **Create**.

Result

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

Initialize a replication relationship in NetApp Workload Factory

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

About this task

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

Before you begin

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

Steps

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

Protect your data with NetApp Autonomous Ransomware Protection with AI

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

About this task

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

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

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

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

Learning and active modes

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

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

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

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

Unsupported configurations

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

- iSCSI volumes
- NVMe volumes

Enable ARP/AI for a file system or a volume

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

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

Before you begin

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

Enable ARP/AI for a file system

Steps

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

Enable ARP/AI for a volume

Steps

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

Validate ransomware attacks

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

Steps

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

Recover data after a ransomware attack

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

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

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

Steps

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

Result

The compromised data has been restored.

Clone a volume in NetApp Workload Factory

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

The clone reflects the current, point-in-time state of the data. You can also use clones to give additional users

access to data without giving them access to production data.

About this task

Volume cloning is only supported for FlexClone volumes.

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

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

Steps

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

Use on-premises ONTAP cluster data in NetApp Workload Factory

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

About this task

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

Use cases

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

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


Discover an on-premises ONTAP cluster

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

Before you begin

Make sure you have the following before you begin:

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



Steps

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

Result

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

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

Replicate volume data from an on-premises ONTAP cluster

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

Before you begin

- You must discover an on-premises ONTAP cluster to replicate its volume data.
- You must have an available FSx for ONTAP file system to be the target for the replication.
- Both the on-premises ONTAP cluster and the FSx for ONTAP file system you use for the replication relationship must have an associated link. [Learn how to associate an existing link or to create and associate a new link](#). After the link associates, return to this operation.

Steps

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

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

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

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

Unlimited.

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

7. Under Replication settings, provide the following:

- a. **Replication interval:** Select the frequency that snapshots are transferred from the source volume to the target volume.
- b. **Long-term retention:** Optionally, enable snapshots for long-term retention.

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

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

8. Select **Create**.

Result

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

Remove an on-premises ONTAP cluster from NetApp Workload Factory

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

Before you begin

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

Steps

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

Result

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

Protect your data with a cyber vault

A cyber vault volume is an isolated, secure storage location used to store backup copies of your data, protecting it from ransomware attacks and other cyber threats. As part of

vault creation, you'll create a cyber vault volume, disable all client protocols, and set up a replication relationship between the source volume and the cyber vault volume, and create immutable snapshots on the cyber vault volume.

What is a cyber vault?

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

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

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

FSx for ONTAP cyber vaults

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

Implementation

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

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

Related information

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

Copyright information

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.