

FSx for ONTAP



FSx for ONTAP: ONTAP User Guide

Table of Contents

What is Amazon FSx for NetApp ONTAP?	1
Features of FSx for ONTAP	2
Security and data protection	3
Pricing for FSx for ONTAP	4
FSx for ONTAP forums	4
Are you a first-time Amazon FSx user?	4
How it works	6
File systems	6
Storage virtual machines	6
Volumes	7
Storage tiers	7
Data tiering	7
Storage efficiency	8
Accessing your data	8
Managing FSx for ONTAP resources	8
Setting up	10
Sign up for an Amazon Web Services account	10
Secure IAM users	10
Next step	11
Getting started	12
Create your FSx for ONTAP file system	12
Step 2: Mounting your file system	14
Step 3: Clean up resources	18
Accessing your data	20
Supported clients	20
Accessing data from within Amazon	21
Accessing data from the same VPC	22
Accessing data from a different VPC	22
Accessing data from on-premises	27
Accessing NFS, SMB, or the ONTAP CLI or REST API endpoints from on-premises	28
Accessing inter-cluster endpoints from on-premises	30
Mounting volumes	30
Mounting on Linux clients	32
Mounting on Windows clients	35

Mounting on macOS clients	37
Mounting iSCSI LUNs	39
Mounting iSCSI LUNs to a Linux client	40
Mounting iSCSI LUNs to a Windows client	51
Using FSx for ONTAP with other Amazon services	58
Using WorkSpaces	58
Using Amazon ECS	64
Using VMware Cloud	67
Availability and durability	68
Choosing a file system deployment type	68
Single-AZ deployment type	68
Multi-AZ deployment type	69
Failover process for FSx for ONTAP	70
Testing failover on a file system	71
Network resources	71
Subnets	72
File system elastic network interfaces	72
Managing storage capacity	74
Storage tiers	74
Choosing file system storage capacity	76
How SSD storage is used	76
Recommended SSD capacity utilization	77
Storage efficiency	78
File system storage capacity and IOPS	79
Scaling SSD storage and IOPS	80
Monitoring SSD storage utilization	82
Setting a storage utilization alarm	83
Viewing storage efficiency savings	86
Modifying SSD storage and IOPS	88
Monitoring storage capacity and IOPS updates	92
Increasing storage capacity dynamically	95
Volume storage capacity	100
Volume data tiering	101
Snapshots and storage capacity	104
Volume file capacity	105
Updating a volume's storage capacity	106

Enabling volume autosizing	107
Monitor volume storage capacity	108
Setting a volume's tiering policy	111
Setting cooling days	113
Setting cloud retrieval policy	116
Viewing a volume's file capacity	117
Increasing the maximum number of files on a volume	117
Enabling cloud write mode	119
Protecting your data	121
Working with backups	121
How backups work	122
Storage requirements	123
Automatic daily backups	123
User-initiated backups	124
Copying tags to backups	124
Backup performance	124
Using Amazon Backup with Amazon FSx	125
Restoring backups to a new volume	126
Deleting backups	127
Backups and offline volumes	127
Creating a user-initiated backup	128
Restoring a backup to a new volume	128
Deleting a backup	131
Working with snapshots	132
Snapshot policies	133
Restoring individual files and folders	134
Restore files from Snapshots	134
Deleting snapshots	135
Create a Snapshot autodelete policy	135
Delete snapshots	136
Disabling automatic snapshots	136
Snapshot reserve	138
Updating Snapshot reserve	139
Scheduled replication	140
Using NetApp BlueXP to schedule replication	140
Using the NetApp ONTAP CLI to schedule replication	141

Protecting data with SnapLock	141
How SnapLock works	141
SnapLock Compliance	146
SnapLock Enterprise	148
Retention period	151
Committing files to WORM	154
Backing up SnapLock volumes	159
Deleting SnapLock volumes	159
Working with Active Directory	161
Self-managed Active Directory prerequisites	162
Self-managed Active Directory requirements	162
Network configuration requirements	162
Active Directory service account requirements	164
Self-managed AD best practices	166
Delegating permissions to your Amazon FSx service account	166
Keep an AD configuration updated	167
Limit traffic within a VPC with security groups	168
Creating outbound security group rules	168
Joining SVMs to an Active Directory	168
Active Directory information needed	169
Managing SVM Active Directory configurations	170
Join an SVM to Active Directory	171
Update an SVM Active Directory configuration using Amazon console, CLI, API	174
Manage Active Directory configuration with NetApp CLI	175
Performance	181
Measuring performance	181
Latency	181
Throughput and IOPS	181
SMB Multichannel and NFS nconnect support	182
Performance details	182
Impact of deployment type on performance	184
Impact of storage capacity on performance	186
Impact of throughput capacity on performance	186
Example: storage capacity and throughput capacity	191
Administering resources	192
Managing file systems	192

File system resources	193
HA pairs	195
Creating FSx for ONTAP file systems	195
Creating file systems in shared subnets	204
Updating a file system	208
Deleting a file system	211
Viewing file system details	211
File system status	212
Managing SVMs	213
Maximum number of SVMs per file system	213
Creating an SVM	214
Updating an SVM	220
Deleting an SVM	222
Viewing SVM details	223
Managing volumes	224
Volume styles	225
Volume types	227
Volume security style	228
Creating volumes	229
Updating a volume	234
Deleting a volume	236
Viewing a volume	237
Creating an iSCSI LUN	237
Next steps	239
Managing SMB shares	239
File access auditing	241
File access auditing overview	241
Overview of tasks for setting up file access auditing	245
Storage capacity and IOPS	252
Throughput capacity	253
When to modify throughput capacity	254
How concurrent throughput and storage scaling requests are handled	254
How to modify throughput capacity	255
Monitoring throughput capacity changes	256
Maintenance windows	258
Tag your resources	259

Tag basics	260
Tagging your resources	261
Copying tags to backups	262
Tag restrictions	262
Permissions and tagging	263
Managing with NetApp applications	263
Signing up for a NetApp account	264
Using NetApp BlueXP	265
Using the NetApp ONTAP CLI	266
Using the ONTAP REST API	270
Security	271
Data protection	272
Data encryption in FSx for ONTAP	273
Encryption at rest	273
Encrypting data in transit	275
Identity and access management	296
Audience	297
Authenticating with identities	297
Managing access using policies	300
FSx for ONTAP and IAM	302
Identity-based policy examples	309
Troubleshooting	312
Using tags with Amazon FSx	314
Using service-linked roles	320
Amazon managed policies	325
AmazonFSxServiceRolePolicy	326
AmazonFSxDeleteServiceLinkedRoleAccess	326
AmazonFSxFullAccess	326
AmazonFSxConsoleFullAccess	327
AmazonFSxConsoleReadOnlyAccess	328
AmazonFSxReadOnlyAccess	329
Policy updates	329
File System Access Control with Amazon VPC	338
Amazon VPC security groups	339
Compliance Validation	341
Interface VPC endpoints	342

Considerations for Amazon FSx interface VPC endpoints	343
Creating an interface VPC endpoint for Amazon FSx API	343
Creating a VPC endpoint policy for Amazon FSx	344
Resilience	344
Backup and restore	344
Snapshots	345
Availability Zones	345
Infrastructure Security	346
Using antivirus software	346
ONTAP roles and users	346
Predefined roles on an SVM	347
Creating new roles or users	350
Updating the fsxadmin account password fails	351
Creating new roles for an SVM using the NetApp ONTAP CLI	353
Using Active Directory user accounts with your file system	354
Configuring public key authentication	356
Migrating to Amazon FSx	358
Migrating using SnapMirror	358
Before you begin	360
Create the destination volume	361
Record the source and destination inter-cluster LIFs	362
Establish cluster peering between source and destination	363
Create an SVM peering relationship	364
Create the SnapMirror relationship	365
Transfer data to your FSx for ONTAP file system	365
Cutting over to Amazon FSx	366
Migrating files with Amazon DataSync	368
Prerequisites	369
DataSync migration basic steps	369
Monitoring file systems	370
Monitoring with CloudWatch	371
How to use FSx for ONTAP CloudWatch metrics	372
Accessing CloudWatch metrics	377
File system metrics	380
Scale-out file system metrics	400
Volume metrics	416

Performance warnings and recommendations	424
Creating alarms	427
Monitoring workload balance	429
Primary storage utilization balance	429
File server and disk performance utilization imbalance	430
Mapping CloudWatch dimensions to ONTAP CLI and REST API resources	431
Rebalancing high-traffic clients	432
Rebalancing highly-utilized volumes	433
Monitoring EMS events	436
Overview of EMS events	436
Viewing EMS events	437
EMS event forwarding to a Syslog server	444
Monitoring with Cloud Insights	446
Monitoring with Harvest and Grafana	447
Getting started with Harvest and Grafana	447
Supported Harvest dashboards	447
Amazon CloudFormation template	447
Amazon EC2 instance types	448
Deployment procedure	449
Logging in to Grafana	452
Troubleshooting Harvest and Grafana	452
Logging with Amazon CloudTrail	455
Amazon FSx Information in CloudTrail	455
Understanding Amazon FSx Log File Entries	456
Quotas	459
Quotas that you can increase	459
Resource quotas for each file system	460
Troubleshooting	464
My Multi-AZ file system is in a MISCONFIGURED state	464
The VPC owner account has disabled Multi-AZ VPC sharing	464
You can't create a new SVM on a Multi-AZ file system	465
You can't access your file system	465
The file system's elastic network interface was modified or deleted	466
The Elastic IP address attached to the file system's elastic network interface was deleted	466
The file system's VPC security group lacks the required inbound rules	466

The compute instance's VPC security group lacks the required outbound rules	466
The compute instance's subnet doesn't use any of the route tables associated with your file system	466
Amazon FSx can't update route table for Multi-AZ file systems created using Amazon CloudFormation	467
Can't access a file system over iSCSI from a client in another VPC	467
The owning account has unshared the VPC subnet	467
Can't access a file system over NFS, SMB, the ONTAP CLI, or the ONTAP REST API from a client in another VPC or on-premises	468
You can't join a storage virtual machine (SVM) to Active Directory	468
The SVM NetBIOS name is the same as the NetBIOS name for the home domain.	469
The SVM is already joined to another Active Directory	469
Amazon FSx can't connect to your Active Directory domain controllers because the SVM's NetBIOS name is already in use	470
Amazon FSx can't communicate with your Active Directory domain controllers	470
Amazon FSx can't connect to your Active Directory due to unmet port requirements or service account permissions	470
Amazon FSx can't connect to your Active Directory domain controllers because the service account credentials are not valid	471
Amazon FSx can't connect to your Active Directory domain controllers because of insufficient service account credentials	472
Amazon FSx can't communicate with your Active Directory DNS servers or domain controllers	472
Amazon FSx can't communicate with your Active Directory because of a invalid Active Directory domain name.	475
The service account can't access the administrators group specified in the SVM Active Directory configuration	475
Amazon FSx can't connect to the Active Directory domain controllers because the organizational unit specified doesn't exist or isn't accessible	476
You can't delete a storage virtual machine or volume	476
Identifying failed deletions	477
SVM deletion: Route tables inaccessible	478
SVM deletion: Peer relationship	480
SVM or volume deletion: SnapMirror	481
SVM deletion: Kerberos-enabled LIF	482
SVM deletion: Other reason	484

Volume deletion: FlexCache relationship	486
Automatic daily backups fail due to insufficient volume capacity	487
You have insufficient volume capacity	487
Determine how your volume storage capacity is being used	487
Increasing a volume's storage capacity	488
Using volume autosizing	488
Your file system's primary storage is full	488
Deleting snapshots	488
Increasing a volume's maximum file capacity	489
Troubleshooting network issues	489
You want to capture a packet trace	489
Document history	493

What is Amazon FSx for NetApp ONTAP?

Amazon FSx for NetApp ONTAP is a fully managed service that provides highly reliable, scalable, high-performing, and feature-rich file storage built on NetApp's popular ONTAP file system. FSx for ONTAP combines the familiar features, performance, capabilities, and API operations of NetApp file systems with the agility, scalability, and simplicity of a fully managed Amazon Web Service.

FSx for ONTAP provides feature-rich, fast, and flexible shared file storage that's broadly accessible from Linux, Windows, and macOS compute instances running in Amazon or on premises. FSx for ONTAP offers high-performance solid state drive (SSD) storage with submillisecond latencies. With FSx for ONTAP, you can achieve SSD levels of performance for your workload while paying for SSD storage for only a small fraction of your data.

Managing your data with FSx for ONTAP is easier because you can snapshot, clone, and replicate your files with the click of a button. In addition, FSx for ONTAP automatically tiers your data to lower-cost, elastic storage, lessening the need for you to provision or manage capacity.

FSx for ONTAP also provides highly available and durable storage with fully managed backups and support for cross-Region disaster recovery. To make it easier to protect and secure your data, FSx for ONTAP supports popular data security and antivirus applications.

For customers who use NetApp ONTAP on-premises, FSx for ONTAP is an ideal solution to migrate, back up, or burst your file-based applications from on-premises to Amazon without the need to change your application code or how you manage your data.

As a fully managed service, FSx for ONTAP makes it easier to launch and scale reliable, high-performing, and secure shared file storage in the cloud. With FSx for ONTAP, you no longer have to worry about:

- Setting up and provisioning file servers and storage volumes
- Replicating data
- Installing and patching file server software
- Detecting and addressing hardware failures
- Managing failover and failback
- Manually performing backups

FSx for ONTAP also provides rich integration with other Amazon services, such as Amazon Identity and Access Management (IAM), Amazon WorkSpaces, Amazon Key Management Service (Amazon KMS), and Amazon CloudTrail.

Topics

- [Features of FSx for ONTAP](#)
- [Security and data protection](#)
- [Pricing for FSx for ONTAP](#)
- [FSx for ONTAP forums](#)
- [Are you a first-time Amazon FSx user?](#)

Features of FSx for ONTAP

With FSx for ONTAP, you get a fully managed file storage solution with:

- Support for petabyte-scale datasets in a single namespace
- Up to tens of gigabytes per second (GBps) of throughput per file system
- Multi-protocol access to data using the Network File System (NFS), Server Message Block (SMB), and Internet Small Computer Systems Interface (iSCSI) protocols
- Highly available and durable Multi-AZ and Single-AZ deployment options
- Automatic data-tiering that reduces storage costs by automatically transitioning infrequently accessed data to a lower-cost storage tier based on your access patterns
- Data compression, deduplication, and compaction to reduce your storage consumption
- Support for NetApp's SnapMirror replication feature
- Support for NetApp's on-premises caching solutions: NetApp Global File Cache and FlexCache
- Support for access and management using native Amazon or NetApp tools and API operations
 - Amazon Web Services Management Console, Amazon Command Line Interface (Amazon CLI), and SDKs
 - NetApp ONTAP CLI, REST API, and BlueXP
- Support for the following data protection and security features:
 - Encryption of file system data and backups at rest using Amazon KMS keys
 - Encryption of data in transit using SMB Kerberos session keys
 - On-demand antivirus scanning

- Authentication and authorization using Microsoft Active Directory
- File access auditing
- NetApp's SnapLock feature with support for both Compliance and Enterprise volumes

Security and data protection

Amazon FSx provides multiple levels of security and compliance to facilitate protecting your data. It automatically encrypts data at rest in file systems and backups using keys that you manage in Amazon Key Management Service (Amazon KMS). You can also encrypt data in transit using Kerberos for NFS and SMB clients.

Amazon FSx has been assessed to comply with the following standards:

- International Standards Organization (ISO)
- Payment Card Industry Data Security Standard (PCI DSS)
- System and Organization Controls (SOC) certifications
- The Health Insurance Portability and Accountability Act of 1996 (HIPAA)

For more information, see [Data protection in Amazon FSx for NetApp ONTAP](#).

Amazon FSx also provides the following levels of access control:

- At the file system level, Amazon FSx provides access control by using Amazon Virtual Private Cloud (Amazon VPC) security groups.
- At the API level, Amazon FSx provides access control by using Amazon Identity and Access Management (IAM) access policies.
- To provide access control at the file and folder level, Amazon FSx supports Unix permissions, NFS access control lists (ACLs), and NTFS ACLs. When you join Amazon FSx to an Active Directory, users who are accessing file systems can authenticate using their Active Directory credentials.

So that you can see the actions taken by users on your Amazon FSx resources, Amazon FSx integrates with Amazon CloudTrail to monitor and log your Amazon FSx API calls. For more information, see [Logging FSx for ONTAP API Calls with Amazon CloudTrail](#).

Additionally, Amazon FSx protects your data with highly durable file system backups. Amazon FSx performs automatic daily backups, and you can take additional backups at any point. For more information, see [Protecting your data](#).

Pricing for FSx for ONTAP

You are billed for file systems based on the following categories:

- SSD storage capacity (per gigabyte-month, or GB-month)
- SSD IOPS that you provision above three IOPS/GB (per IOPS-month)
- Throughput capacity (per megabytes per second [MBps]-month)
- Capacity pool storage consumption (per GB-month)
- Capacity pool requests (per read and write)
- Backup storage consumption (per GB-month)

For more information about pricing and fees associated with the service, see [Amazon FSx for NetApp ONTAP pricing](#).

FSx for ONTAP forums

If you encounter issues while using Amazon FSx, use the FSx for ONTAP discussion [forums](#) to get answers.

Are you a first-time Amazon FSx user?

If you're a first-time user of Amazon FSx, we recommend that you read the following sections in order:

1. If you're new to Amazon, see [Setting up FSx for ONTAP](#) to set up an Amazon Web Services account.
2. If you're ready to create your first Amazon FSx file system, follow the instructions in [Getting started with Amazon FSx for NetApp ONTAP](#).
3. For information about performance, see [Amazon FSx for NetApp ONTAP performance](#).
4. For Amazon FSx security details, see [Security in Amazon FSx for NetApp ONTAP](#).

5. For information about the Amazon FSx API, see the [Amazon FSx API Reference](#).

How Amazon FSx for NetApp ONTAP works

This topic introduces the major features of Amazon FSx for NetApp ONTAP file systems and how they work, with links to sections with in-depth descriptions, important implementation details, and step-by-step configuration procedures.

Topics

- [FSx for ONTAP file systems](#)
- [Storage virtual machines](#)
- [Volumes](#)
- [Storage tiers](#)
- [Storage efficiency](#)
- [Accessing data stored on FSx for ONTAP file systems](#)
- [Managing FSx for ONTAP resources](#)

FSx for ONTAP file systems

A file system is the primary FSx for ONTAP resource, analogous to an on-premises NetApp ONTAP cluster. You specify the solid state drive (SSD) storage capacity and throughput capacity for your file system, and choose an Amazon Virtual Private Cloud (VPC) where your file system is created. For more information, see [Managing FSx for ONTAP file systems](#).

Your file system can have one to 12 high-availability (HA) pairs depending on its configuration. An HA pair is made up of two file servers in an active-standby configuration. File systems with a single HA pair are called *scale-up* file systems. File systems with multiple HA pairs are called *scale-out* file systems. For more information, see [High-availability \(HA\) pairs](#).

Storage virtual machines

A storage virtual machine (SVM) is an isolated file server with its own administrative and data access endpoints for administering and accessing data. When you access data in your FSx for ONTAP file system, your clients and workstations interface with an SVM using the SVM's endpoint IP address. For more information, see [Managing SVMs](#).

You can join SVMs to a Microsoft Active Directory for file access authentication and authorization. For more information, see [Working with Microsoft Active Directory in FSx for ONTAP](#).

Volumes

FSx for ONTAP **volumes** are virtual resources that you use for organizing and grouping your data. Volumes are logical containers that are hosted on SVMs, and data stored in them consumes physical storage capacity on your file system.

When you create a volume, you set its size, which determines the amount of physical data that you can store in it, regardless of which storage tier the data is stored on. You also set the volume type, either RW (read-writable) or DP (data protection). A DP volume is read-only and can be used as the destination in a NetApp SnapMirror or SnapVault relationship.

FSx for ONTAP volumes are thin provisioned, meaning that they only consume storage capacity for the data stored in them. With thin-provisioned volumes, storage capacity is not reserved in advance. Instead, storage is allocated dynamically, as it is needed. Free space is released back to the file system when data in the volume or LUN is deleted. For example, you can create three 10 TiB volumes on a file system configured with 10 TiB of free storage capacity, as long as the total amount of data stored in the three volumes doesn't exceed 10 TiB at any time. The amount of data physically stored on a volume counts toward your overall storage capacity consumption. For more information, see [Managing FSx for ONTAP volumes](#).

Storage tiers

An FSx for ONTAP file system has two *storage tiers*: primary storage and capacity pool storage. Primary storage is provisioned, scalable, high-performance SSD storage that's purpose-built for the active portion of your data set. Capacity pool storage is a fully elastic storage tier that can scale to petabytes in size and is cost optimized for infrequently accessed data. Data that you write to your volumes consumes capacity on your storage tiers. For more information, see [FSx for ONTAP storage tiers](#).

Data tiering

Data tiering is the process by which Amazon FSx for NetApp ONTAP automatically moves data between the *SSD* and the *capacity pool* storage tiers. Each volume has a tiering policy that controls whether data is moved to the capacity tier when it becomes inactive (cold). A volume's Tiering

policy cooling period determines when data becomes inactive (cold). For more information, see [Volume data tiering](#).

Storage efficiency

Amazon FSx for NetApp ONTAP supports ONTAP's block-level storage efficiency features—compaction, compression, and deduplication—to reduce the storage capacity that your data consumes. Storage efficiency features can reduce the footprint of your data in SSD storage, capacity pool storage, and backups. The typical storage capacity savings for general purpose file sharing workloads without sacrificing performance is 65% from compression, deduplication, and compaction, on both the SSD and capacity pool storage tiers. For more information, see [FSx for ONTAP storage efficiency](#).

Accessing data stored on FSx for ONTAP file systems

You can access your data on FSx for ONTAP volumes from multiple Linux, Windows, or macOS clients simultaneously over the NFS (v3, v4, v4.1, v4.2) and SMB protocols. You can also access data using the iSCSI (block) protocol. For more information, see [Accessing data](#).

Managing FSx for ONTAP resources

There are several ways that you can interact with your FSx for ONTAP file system and manage its resources. You can manage your FSx for ONTAP resources using both Amazon and NetApp ONTAP management tools:

- Amazon management tools
 - The Amazon Web Services Management Console
 - The Amazon Command Line Interface (Amazon CLI)
 - The Amazon FSx API and SDKs
 - Amazon CloudFormation
- NetApp management tools:
 - NetApp BlueXP
 - The NetApp ONTAP CLI
 - The NetApp ONTAP REST API

For more information, see [Administering resources](#).

Setting up FSx for ONTAP

Before you use Amazon FSx for the first time, complete the following tasks:

1. [Sign up for an Amazon Web Services account](#)
2. [Secure IAM users](#)

Topics

- [Sign up for an Amazon Web Services account](#)
- [Secure IAM users](#)
- [Next step](#)

Sign up for an Amazon Web Services account

If you do not have an Amazon Web Services account, use the following procedure to create one.

To sign up for Amazon Web Services

1. Open <http://www.amazonaws.cn/> and choose **Sign Up**.
2. Follow the on-screen instructions.

Amazon sends you a confirmation email after the sign-up process is complete. At any time, you can view your current account activity and manage your account by going to <http://www.amazonaws.cn/> and choosing **My Account**.

Secure IAM users

After you sign up for an Amazon Web Services account, safeguard your administrative user by turning on multi-factor authentication (MFA). For instructions, see [Enable a virtual MFA device for an IAM user \(console\)](#) in the *IAM User Guide*.

To give other users access to your Amazon Web Services account resources, create IAM users. To secure your IAM users, turn on MFA and only give the IAM users the permissions needed to perform their tasks.

For more information about creating and securing IAM users, see the following topics in the *IAM User Guide*:

- [Creating an IAM user in your Amazon Web Services account](#)
- [Access management for Amazon resources](#)
- [Example IAM identity-based policies](#)

Next step

To get started using FSx for ONTAP see [Getting started with Amazon FSx for NetApp ONTAP](#) for instructions to create your Amazon FSx resources.

Getting started with Amazon FSx for NetApp ONTAP

Learn how to get started using Amazon FSx for NetApp ONTAP. This getting started exercise includes the following steps.

Topics

- [Step 1: Create an Amazon FSx for NetApp ONTAP file system](#)
- [Step 2: Mounting your file system from an Amazon EC2 Linux instance](#)
- [Step 3: Clean up resources](#)

Step 1: Create an Amazon FSx for NetApp ONTAP file system

The Amazon FSx console has two options for creating a file system – a **Quick create** option and a **Standard create** option. To rapidly and easily create an Amazon FSx for NetApp ONTAP file system with the service recommended configuration, use the **Quick create** option.

The **Quick create** option creates a file system with a single high-availability pair (HA), single storage virtual machine (SVM) and a single volume. The **Quick create** option configures this file system to allow data access from Linux instances over the Network File System (NFS) protocol. After your file system is created, you can create additional SVMs and volumes as needed, including an SVM joined to an Active Directory to allow access from Windows and macOS clients over the Server Message Block (SMB) protocol.

For information about using the **Standard create** option to create a file system with a customized configuration, and for using the Amazon CLI and API, see [Creating FSx for ONTAP file systems](#).

To create your file system

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. On the dashboard, choose **Create file system** to start the file system creation wizard.
3. On the **Select file system type** page, choose **Amazon FSx for NetApp ONTAP**, and then choose **Next**. The **Create ONTAP file system** page appears.
4. For **Creation method**, choose **Quick create**.
5. In the **Quick configuration** section, for **File system name - optional**, enter a name for your file system. It's easier to find and manage your file systems when you name them. You can use a

maximum of 256 Unicode letters, white space, and numbers, plus these special characters: + - (hyphen) = . _ (underscore) : /

6. For **Deployment type** choose **Multi-AZ** or **Single-AZ**.
 - **Multi-AZ** file systems replicate your data and support failover across multiple Availability Zones in the same Amazon Web Services Region.
 - **Single-AZ** file systems replicate your data and offer automatic failover within a single Availability Zone.

For more information, see [Availability and durability](#).

7. For **SSD storage capacity**, specify the storage capacity of your file system, in gibibytes (GiBs). Enter any whole number in the range of 1,024–196,608. If you need more SSD storage capacity, you can use **Standard create**. For more information, see [To create a file system \(console\)](#).

You can increase the amount of storage capacity as needed at any time after you create the file system. For more information, see [Managing storage capacity](#).

8. For **Throughput capacity**, Amazon FSx automatically provides a recommended throughput capacity based on your SSD storage. You can also choose your file system's throughput (up to 4,096 MBps). If you need more throughput capacity, you can use **Standard create**.
9. For **Virtual Private Cloud (VPC)**, choose the Amazon VPC that you want to associate with your file system.
10. For **Storage efficiency**, choose **Enabled** to turn on the ONTAP storage efficiency features (compression, deduplication, and compaction) or **Disabled** to turn them off.
11. (Multi-AZ only) **Endpoint IP address range** specifies the IP address range in which the endpoints to access your file system are created.

Choose a **Quick create** option for the endpoint IP address range:

- **Unallocated IP address range from your VPC** – Choose this option to have Amazon FSx use the last 64 IP addresses from the VPC's primary CIDR range as the endpoint IP address range for the file system. Note that this range is shared across multiple file systems if you choose this option multiple times.

Note

- Each file system that you create consumes two IP addresses from this range—one for the cluster, and one for the first SVM. The first and last IP addresses are also reserved. For every additional SVM, the file system consumes another IP address. For example, a file system that hosts 10 SVMs uses 11 IP addresses. Additional file systems work in the same way. They consume the two initial IP addresses, plus one for each additional SVM. The maximum number of file systems using the same IP address range, each with a single SVM, is 31.
 - This option is grayed out if any of the last 64 IP addresses in a VPC's primary CIDR range are in use by a subnet.
- **Floating IP address range outside your VPC** – Choose this option to have Amazon FSx use a 198.19.x.0/24 address range that isn't already used by any other file systems with the same VPC and route tables.

You can also specify your own IP address range in the **Standard create** option.

12. Choose **Next**, and review the file system configuration on the **Create ONTAP file system** page. Note which file system settings you can modify after the file system is created.
13. Choose **Create file system**.

Quick create creates a file system with one SVM (named `fsx`) and one volume (named `vol1`). The volume has a junction path of `/vol1` and a capacity pool tiering policy of **Auto** (which will automatically tier any data that hasn't been accessed for 31 days to lower-cost capacity pool storage). The default snapshot policy gets assigned to the default volume. The file system data is encrypted at rest using your default service managed Amazon KMS key.

Step 2: Mounting your file system from an Amazon EC2 Linux instance

You can mount your file system from an Amazon Elastic Compute Cloud (Amazon EC2) instance. This procedure uses an instance running Amazon Linux 2.







To mount your file system from Amazon EC2

1. Open the Amazon EC2 console at <https://console.amazonaws.cn/ec2/>.
2. Create or select an Amazon EC2 instance running Amazon Linux 2 that is in the same virtual private cloud (VPC) as your file system. For more information about launching an instance, see [Step 1: Launch an instance](#) in the *Amazon EC2 User Guide for Linux Instances*.
3. Connect to your Amazon EC2 Linux instance. For more information, see [Connect to your Linux instance](#) in the *Amazon EC2 User Guide for Linux Instances*.
4. Open a terminal on your Amazon EC2 instance using secure shell (SSH), and log in with the appropriate credentials.
5. Create a directory on your Amazon EC2 instance to use as the volume's mount point with the following command. In the following example, replace *mount-point* with your own information.

```
$ sudo mkdir /mount-point
```

6. Mount your Amazon FSx for NetApp ONTAP file system to the directory that you created. Use a mount command similar to the example that follows. In the following example, replace the following placeholder values with your own information.
 - *nfs_version* – The NFS version you are using; FSx for ONTAP supports versions 3, 4.0, 4.1, and 4.2.
 - *nfs-dns-name* – The NFS DNS name of the storage virtual machine (SVM) in which the volume you are mounting exists. You can find the NFS DNS name in the Amazon FSx console by choosing **Storage virtual machines**, then choosing the SVM on which the volume you are mounting exists. The NFS DNS name is found on the **Endpoints** panel, shown in the following image.

Endpoints

Management DNS name svm-0123456789abcdefa.fs- 0123456789abcdefa.fsx.us-east-2.amazonaws.com 	Management IP address 198.51.100.1 
NFS DNS name svm-0123456789abcdefa.fs- 0123456789abcdefa.fsx.us-east-2.amazonaws.com 	NFS IP address 198.51.100.1 
iSCSI DNS name iscsi-svm-0123456789abcdefa.fs- 0123456789abcdefa.fsx.us-east-2.amazonaws.com 	iSCSI IP addresses 198.51.100.37,198.51.100.123 

- *volume-junction-path* – The junction path of the volume that you're mounting. You can find a volume's junction path in the Amazon FSx console on the **Summary** panel of the Volume details page, shown in the following image.

vol1 (fsvol-0123456789abcdef2)

Attach

Actions ▼

Summary

Volume ID

fsvol-0123456789abcdef2 

Creation time

2022-09-06T15:02:38-04:00

SVM ID

[svm-abcdef0123456789f](#)


Volume name

vol1 

Lifecycle state

 Created

Junction path

/vol1 

UUID

2248c29a-2e1a-11ed-888b-
a96e652919ea

Volume type

ONTAP


Tiering policy name

AUTO

File system ID


[fs-0468008f689bebaa3](#) 

Size

1.00 TB Tiering policy cooling period
(days)

31

Resource ARN

arn:aws:fsx:us-east-
2:267731178466:volume/fs-
0468008f689bebaa3/fsvol-
0123456789abcdef2 

Storage efficiency enabled

Disabled

- *mount-point* – The name of the directory that you created on your EC2 instance for the volume's mount point.

```
sudo mount -t nfs -o nfsvers=nfs_version nfs-dns-name:/volume-junction-path /mount-point
```

The following command uses example values.

```
sudo mount -t nfs -o nfsvers=4.1 svm-abcdef1234567890c.fs-012345abcdef6789b.fsx.us-east-2.amazonaws.com:/vol1 /fsxN
```

If you have issues with your Amazon EC2 instance (such as connections timing out), see [Troubleshoot EC2 instances](#) in the *Amazon EC2 User Guide for Linux Instances*.

Step 3: Clean up resources

After you have finished this exercise, you should follow these steps to clean up your resources and protect your Amazon Web Services account.

To clean up resources

1. On the Amazon EC2 console, terminate your instance. For more information, see [Terminate Your Instance](#) in the *Amazon EC2 User Guide for Linux Instances*.
2. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
3. On the Amazon FSx console, delete all of your FSx for ONTAP volumes that are not root volumes of your SVM. For more information, see [Deleting a volume](#).
4. Delete all of your FSx for ONTAP SVMs. For more information, see [Deleting a storage virtual machine \(SVM\)](#).
5. On the Amazon FSx console, delete your file system. When you delete a file system, all automatic backups are deleted automatically. However, you still must delete any manually created backups. The following steps outline this process.
 - a. From the console dashboard, choose the name of the file system that you created for this exercise.
 - b. For **Actions**, choose **Delete file system**.
 - c. In the **Delete file system** dialog box, enter the ID of the file system that you want to delete in the **File system ID** box.
 - d. Choose **Delete file system**.
 - e. While Amazon FSx deletes the file system, its status in the dashboard changes to **DELETING**. Once the file system is deleted, it no longer appears in the dashboard. Any automatic backups are deleted along with the file system.
 - f. Now you can delete any manually created backups for your file system. From the left-side navigation, choose **Backups**.
 - g. From the dashboard, choose any backups that have the same **File system ID** as the file system that you deleted, and choose **Delete backup**. Be sure to retain the final backup, if you created one.
 - h. The **Delete backups** dialog box opens. Keep the check box selected for the IDs of the backups that you want to delete, and then choose **Delete backups**.

Your Amazon FSx file system and any related automatic backups are now deleted, along with any manual backups that you chose to delete as well.

Accessing data

You can access your Amazon FSx file systems using a variety of supported clients and methods in both the Amazon Web Services Cloud and on premises environments.

Each SVM has four endpoints that are used to access data or to manage the SVM using the NetApp ONTAP CLI or REST API:

- `Nfs` – For connecting using the Network File System (NFS) protocol
- `Smb` – For connecting using the Service Message Block (SMB) protocol (If your SVM is joined to an Active Directory, or you're using a workgroup.)
- `Iscsi` – For connecting using the Internet Small Computer Systems Interface (iSCSI) protocol (for scale-up file systems only).
- `Management` – For managing SVMs using the NetApp ONTAP CLI or API, or NetApp BlueXP

Topics

- [Supported clients](#)
- [Accessing data from within Amazon](#)
- [Accessing data from on-premises](#)
- [Mounting volumes](#)
- [Mounting iSCSI LUNs](#)
- [Using FSx for ONTAP with other Amazon services](#)

Supported clients

FSx for ONTAP file systems support accessing data from a wide variety of compute instances and operating systems. It does this by supporting access using the Network File System (NFS) protocol (v3, v4.0, v4.1 and v4.2), all versions of the Server Message Block (SMB) protocol (including 2.0, 3.0, and 3.1.1), and the Internet Small Computer Systems Interface (iSCSI) protocol.

⚠ Important

Amazon FSx doesn't support accessing file systems from the public internet. Amazon FSx automatically detaches any Elastic IP address which is a public IP address reachable from the Internet, that gets attached to a file system's elastic network interface.

The following Amazon compute instances are supported for use with FSx for ONTAP:

- Amazon Elastic Compute Cloud (Amazon EC2) instances running Linux with NFS or SMB support, Microsoft Windows, and MacOS. For more information, see [Mounting volumes](#).
- Amazon Elastic Container Service (Amazon ECS) Docker containers on Amazon EC2 Windows and Linux instances. For more information, see [Using Amazon Elastic Container Service with FSx for ONTAP](#).
- Amazon Elastic Kubernetes Service – To learn more, see [Amazon FSx for NetApp ONTAP CSI driver](#) in the Amazon EKS User Guide.
- Red Hat OpenShift Service on Amazon (ROSA) – To learn more, see [What is Red Hat OpenShift Service on Amazon?](#) in the *Red Hat OpenShift Service on Amazon User Guide*.
- Amazon WorkSpaces instances. For more information, see [Using Amazon WorkSpaces with FSx for ONTAP](#).
- Amazon AppStream 2.0 instances.
- Amazon Lambda – For more information, see the Amazon blog post [Enabling SMB access for server-less workloads with Amazon FSx](#).
- Virtual machines (VMs) running in VMware Cloud on Amazon environments. For more information, see [Configure Amazon FSx for NetApp ONTAP as External Storage](#) and [VMware Cloud on Amazon with Amazon FSx for NetApp ONTAP Deployment Guide](#).

Once mounted, FSx for ONTAP file systems appear as a local directory or drive letter over NFS and SMB, providing fully managed, shared network file storage that can be simultaneously accessed by up to thousands of clients. iSCSI LUNS are accessible as block devices when mounted over iSCSI.

Accessing data from within Amazon

Each Amazon FSx file system is associated with a Virtual Private Cloud (VPC). You can access your FSx for ONTAP file system from anywhere in the file system's VPC, regardless of Availability Zone.

You can also access your file system from other VPCs that can be in different Amazon accounts or Amazon Web Services Regions. In addition to the requirements described in the following sections for accessing FSx for ONTAP resources, you also need to ensure that your file system's VPC security group is configured so that data and management traffic can flow between your file system and clients. For more information about configuring security groups with the required ports, see [Amazon VPC security groups](#).

Topics

- [Accessing data from within the same VPC](#)
- [Accessing data from outside the deployment VPC](#)

Accessing data from within the same VPC

When you create your Amazon FSx for NetApp ONTAP file system, you select the Amazon VPC in which it is located. All SVMs and volumes associated with the Amazon FSx for NetApp ONTAP file system are also located in the same VPC. When mounting a volume, if the file system and the client mounting the volume are located in the same VPC and Amazon Web Services account, you can use the SVM's DNS name and volume junction or SMB share, depending on the client. For more information, see [Mounting volumes](#).

You can achieve optimal performance if the client and the volume are located in the in the same Availability Zone as the file system's subnet, or preferred subnet for Multi-AZ file systems. To identify a file system's subnet or preferred subnet, in the Amazon FSx console, choose **File systems**, then choose the ONTAP file system whose volume you are mounting, and the subnet or preferred subnet (Multi-AZ) is displayed in the **Subnet** or **Preferred subnet** panel.

Accessing data from outside the deployment VPC

This section describes how to access an FSx for ONTAP file system's endpoints from Amazon locations outside of the file system's deployment VPC.

Accessing NFS, SMB, and ONTAP management endpoints on Multi-AZ file systems


The NFS, SMB, and ONTAP management endpoints on Amazon FSx for NetApp ONTAP Multi-AZ file systems use floating internet protocol (IP) addresses so that connected clients seamlessly transition between the preferred and standby file servers during a failover event. For more information about failovers, see [Failover process for FSx for ONTAP](#).

These floating IP addresses are created in the VPC route tables that you associate with your file system, and are within the file system's `EndpointIpAddressRange` which you can specify during creation. The `EndpointIpAddressRange` uses the following address ranges, depending on how a file system is created:

- Multi-AZ file systems created using the Amazon FSx console use the last 64 IP addresses in the VPC's primary CIDR range for the file system's `EndpointIpAddressRange` by default.
- Multi-AZ file systems created using the Amazon CLI or Amazon FSx API use an IP address range within the `198.19.0.0/16` address block for the `EndpointIpAddressRange` by default.

Only [Amazon Transit Gateway](#) supports routing to floating IP addresses, which is also known as transitive peering. VPC Peering, Amazon Direct Connect, and Amazon VPN don't support transitive peering. Therefore, you are required to use Transit Gateway in order to access these interfaces from networks that are outside of your file system's VPC.

The following diagram illustrates using Transit Gateway for NFS, SMB, or management access to a Multi-AZ file system that is in a different VPC than the clients that are accessing it.



Note

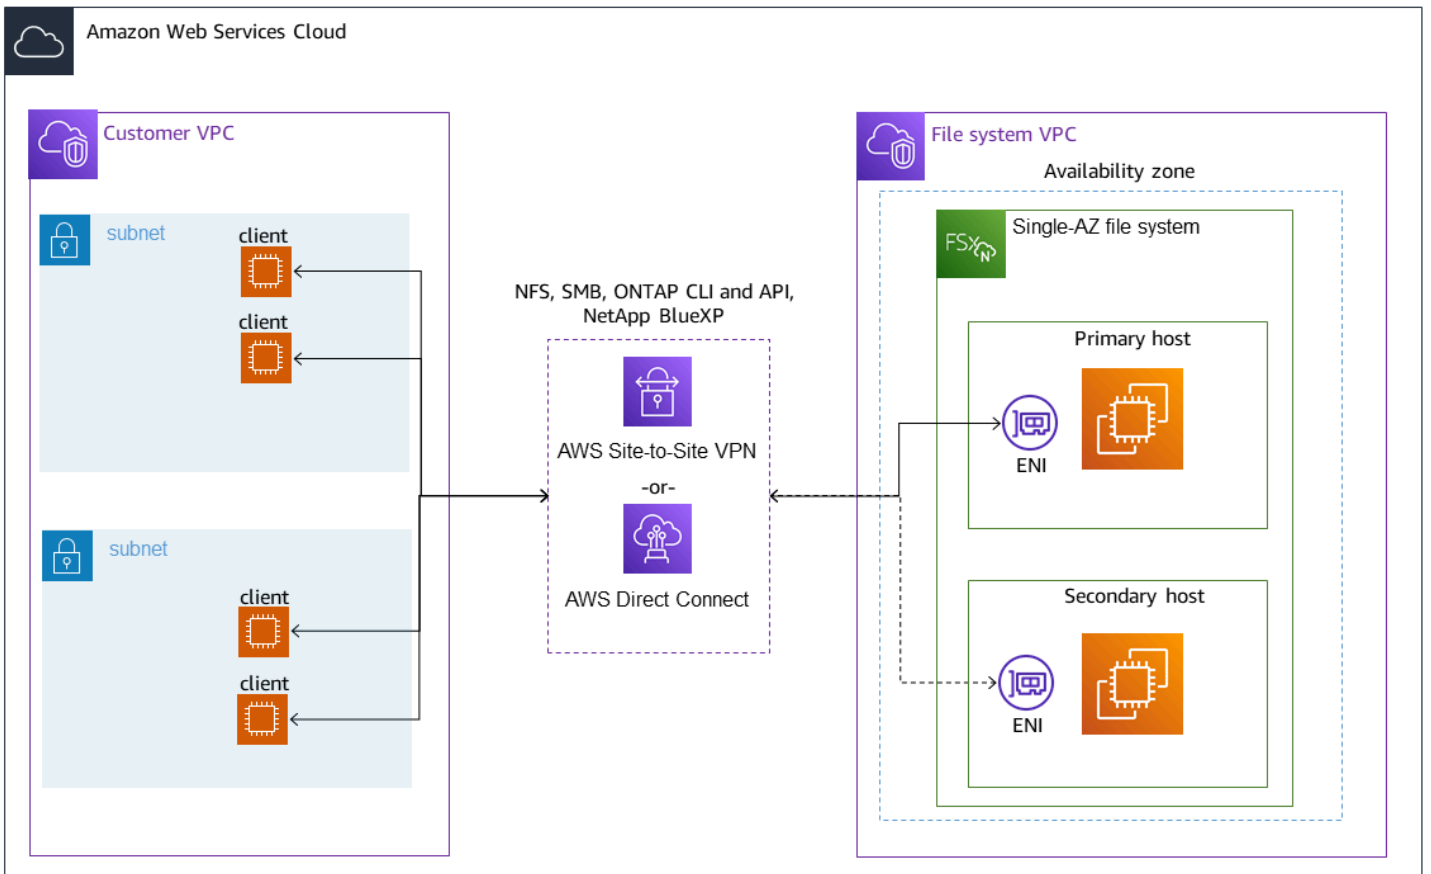
Ensure that all of the route tables you're using are associated with your Multi-AZ file system. Doing so helps prevent unavailability during a failover. For information about associating your Amazon VPC route tables with your file system, see [Updating a file system](#).

For information about when you need to use Transit Gateway to access your FSx for ONTAP file system, see [When is Transit Gateway required?](#).

Accessing NFS, SMB, or the ONTAP CLI and API for Single-AZ file systems

The endpoints used to access FSx for ONTAP Single-AZ file systems over NFS or SMB, and for administering file systems using the ONTAP CLI or REST API, are secondary IP addresses on the ENI of the active file server. The secondary IP addresses are within the VPC's CIDR range, so clients can access data and management ports using VPC Peering, Amazon Direct Connect, or Amazon VPN without requiring Amazon Transit Gateway.

The following diagram illustrates using Amazon VPN or Amazon Direct Connect for NFS, SMB, or management access to a Single-AZ file system that is in a different VPC than the clients accessing it.



When is Transit Gateway required?

Whether or not Transit Gateway is required for your Multi-AZ file systems depends on the method you use to access your file system data. Single-AZ file systems do not require Transit Gateway. The following table describes when you will need to use Amazon Transit Gateway to access Multi-AZ file systems.

Data access	Requires Transit Gateway?
Accessing FSx over NFS, SMB, or the NetApp ONTAP REST API, CLI or BlueXP	Only if: <ul style="list-style-type: none"> • Accessing from a peered (on-premises, for example) network, and • You are not accessing FSx through a NetApp FlexCache or Global File Cache instance
Accessing data over iSCSI	No

Data access	Requires Transit Gateway?
Joining an SVM to an Active Directory	No
SnapMirror	No
FlexCache Caching	No
Global File Cache	No

Configuring routing using Amazon Transit Gateway

If you have a Multi-AZ file system with an `EndpointIPAddressRange` that's outside your VPC's CIDR range, you need to set up additional routing in your Amazon Transit Gateway to access your file system from peered or on-premises networks.

Important

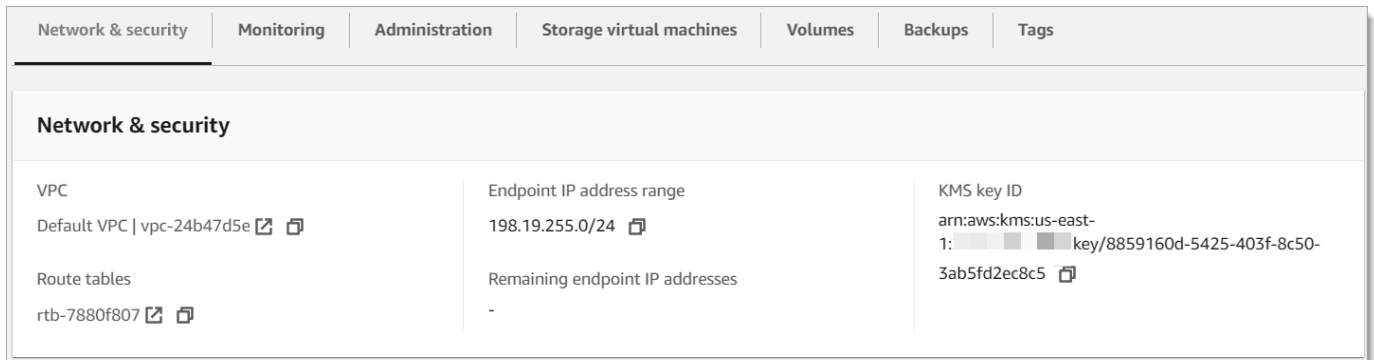
To access a Multi-AZ file system using a Transit Gateway, each of the Transit Gateway's attachments must be created in a subnet whose route table is associated with your file system.

Note

No additional Transit Gateway configuration is required for Single-AZ file systems or Multi-AZ file systems with an `EndpointIPAddressRange` that's within your VPC's IP address range.

To configure routing using Amazon Transit Gateway

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. Choose the FSx for ONTAP file system for which you are configuring access from a peered network.
3. In **Network & security** copy the **Endpoint IP address range**.



- Add a route to Transit Gateway that routes traffic destined for this IP address range to your file system's VPC. For more information, see [Work with transit gateways](#) in the *Amazon VPC Transit Gateways*.
- Confirm that you can access your FSx for ONTAP file system from the peered network.

To add the route table to your file system, see [Updating a file system](#).

Note

DNS records for the management, NFS, and SMB endpoints are only resolvable from within the same VPC as the file system. In order to mount a volume or connect to a management port from another network, you need to use the endpoint's IP address. These IP addresses do not change over time.

Accessing iSCSI or inter-cluster endpoints outside of the deployment VPC

You can use either VPC Peering or Amazon Transit Gateway to access your file system's iSCSI or inter-cluster endpoints from outside of the file system's deployment VPC. You can use VPC Peering to route iSCSI and inter-cluster traffic between VPCs. A VPC peering connection is a networking connection between two VPCs, and is used to route traffic between them using private IPv4 addresses. You can use VPC peering to connect VPCs within the same Amazon Web Services Region or between different Amazon Web Services Regions. For more information on VPC peering, see [What is VPC peering?](#) in the *Amazon VPC Peering Guide*.

Accessing data from on-premises

You can access your FSx for ONTAP file systems from on-premises using [Amazon VPN](#) and [Amazon Direct Connect](#); more specific use case guidelines are available in the following sections. In addition

to any requirements listed below for accessing different FSx for ONTAP resources from on-premises, you also need to ensure that your file system's VPC security group allows data to flow between your file system and clients; for a list of required ports, see [Amazon VPC security groups](#).

Accessing NFS, SMB, or the ONTAP CLI or REST API endpoints from on-premises

This section describes how to access the NFS, SMB, and ONTAP management ports on FSx for ONTAP file systems from on-premises networks.

Accessing Multi-AZ file systems

Amazon FSx requires that you use Amazon Transit Gateway or that you configure remote NetApp Global File Cache or NetApp FlexCache to access Multi-AZ file systems from an on-premises network. In order to support failover across AZs for Multi-AZ file systems, Amazon FSx uses floating IP addresses for the interfaces used for NFS, SMB, and ONTAP management endpoints. Because the NFS, SMB, and management endpoints use floating IPs, you must use [Amazon Transit Gateway](#) in conjunction with Amazon Direct Connect or Amazon VPN to access these interfaces from an on-premises network. The floating IP addresses used for these interfaces are within the `EndpointIpAddressRange` you specify when creating your Multi-AZ file system. If you create your file system from the Amazon FSx console, by default Amazon FSx chooses the last 64 IP addresses from the VPC's primary CIDR range to be used as the endpoint IP address range for the file system. If you create your file system from the Amazon CLI or the Amazon FSx API, by default Amazon FSx chooses an IP address range from within the `198.19.0.0/16` IP address range. The floating IP addresses are used to enable a seamless transition of your clients to the standby file system in the event a failover is required. For more information, see [Failover process for FSx for ONTAP](#).

Important

To access a Multi-AZ file system using a Transit Gateway, each of the Transit Gateway's attachments must be created in a subnet whose route table is associated with your file system.

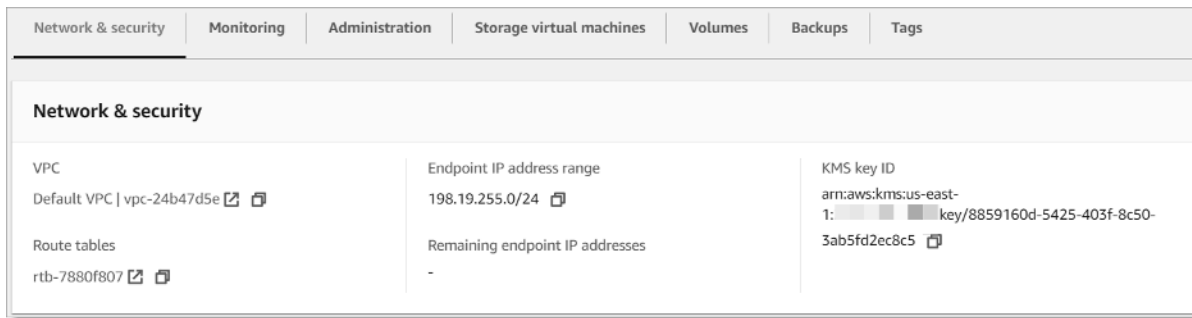
To configure Amazon Transit Gateway for access from outside of your VPC

If you have a Multi-AZ file system with an `EndpointIPAddressRange` that's outside your VPC's CIDR range, you need to set up additional routing in your Amazon Transit Gateway to access your file system from peered or on-premises networks.

Note

No additional Transit Gateway configuration is required for Single-AZ file systems or Multi-AZ file systems with an `EndpointIPAddressRange` that's within your VPC's IP address range.

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. Choose the FSx for ONTAP file system for which you are configuring access from a peered network.
3. In **Network & security** copy the **Endpoint IP address range**.



4. Add a route to the Transit Gateway that routes traffic destined for this IP address range to your file system's VPC. For more information, see [Work with transit gateways](#) in the *Amazon VPC Transit Gateway User Guide*.
5. Confirm that you can access your FSx for ONTAP file system from the peered network.

Important

To access a Multi-AZ file system using a Transit Gateway, each of the Transit Gateway's attachments must be created in a subnet whose route table is associated with your file system.

To add a route table to your file system, see [Updating a file system](#).

Accessing Single-AZ file systems

The requirement to use Amazon Transit Gateway to access data from an on-premises network doesn't exist for Single-AZ file systems. Single-AZ file systems are deployed in a single subnet, and a floating IP address is not required to provide failover between nodes. Instead, the IP addresses you access on Single-AZ file systems are implemented as secondary IP addresses within the file system's VPC CIDR range, enabling you to access your data from another network without requiring Amazon Transit Gateway.







Accessing inter-cluster endpoints from on-premises

FSx for ONTAP's inter-cluster endpoints are dedicated to replication traffic between NetApp ONTAP file systems, including between on-premises NetApp deployments and FSx for ONTAP. Replication traffic includes SnapMirror, FlexCache, and FlexClone relationships between storage virtual machines (SVMs) and volumes across different file systems, and NetApp Global File Cache. The inter-cluster endpoints are also used for Active Directory traffic.

Because a file system's inter-cluster endpoints use IP addresses that are within the CIDR range of the VPC you provide when you create your FSx for ONTAP file system, you are not required to use a Transit Gateway for routing inter-cluster traffic between on-premises and the Amazon Web Services Cloud. However, on-premises clients still must use Amazon VPN or Amazon Direct Connect to establish a secure connection to your VPC.

Mounting volumes

You access the data in FSx for ONTAP by mounting a volume on your client. The commands in this section use the DNS name or the IP address of the SVM in which the volume is created to mount or attach a volume. You can find the SVM's DNS name and IP address in the Amazon FSx console by choosing **ONTAP > Storage virtual machines**, or on the **Storage virtual machine** tab in the **File system details** page for the file system, shown in the following image.

Endpoints	
Management DNS name svm-0123456789abcdefa.fs- 0123456789abcdefa.fsx.us-east-2.amazonaws.com 	Management IP address 198.51.100.1 
NFS DNS name svm-0123456789abcdefa.fs- 0123456789abcdefa.fsx.us-east-2.amazonaws.com 	NFS IP address 198.51.100.1 
iSCSI DNS name iscsi-svm-0123456789abcdefa.fs- 0123456789abcdefa.fsx.us-east-2.amazonaws.com 	iSCSI IP addresses 198.51.100.37,198.51.100.123 

Or, you can find them in the response of the [DescribeStorageVirtualMachines](#) API operation.

You can find a volume's junction path in the Amazon FSx console on the **Summary** panel of the volume details page, shown in the following image.

vol1 (fsvol-0123456789abcdef2)

Attach
Actions ▼

Summary

<p>Volume ID fsvol-0123456789abcdef2 </p> <p>Volume name vol1 </p> <p>UUID 2248c29a-2e1a-11ed-888b-a96e652919ea</p> <p>File system ID fs-0468008f689bebaa3 </p> <p>Resource ARN arn:aws:fsx:us-east-2:267731178466:volume/fs-0468008f689bebaa3/fsvol-0123456789abcdef2 </p>	<p>Creation time 2022-09-06T15:02:38-04:00</p> <p>Lifecycle state Created</p> <p>Volume type ONTAP</p> <p>Size 1.00 TB </p>	<p>SVM ID svm-abcdef0123456789f</p> <p>Junction path /vol1 </p> <p>Tiering policy name AUTO</p> <p>Tiering policy cooling period (days) 31</p> <p>Storage efficiency enabled Disabled</p>
---	--	---

Topics

- [Mounting on Linux clients](#)
- [Mounting on Microsoft Windows clients](#)
- [Mounting on macOS clients](#)

Mounting on Linux clients

We recommend that the SVM volumes to which you are attaching Linux clients have a security style setting of UNIX or mixed. For more information, see [Managing FSx for ONTAP volumes](#).

Note

By default, FSx for ONTAP NFS mounts are hard mounts. To ensure a smooth failover in the event that one occurs, we recommend that you use the default hard mount option.

To mount an ONTAP volume on a Linux client

1. Open the Amazon EC2 console at <https://console.amazonaws.cn/ec2/>.
2. Create or select an Amazon EC2 instance running Amazon Linux 2 that is in the same VPC as the file system.

For more information on launching an EC2 Linux instance, see [Step 1: Launch an instance](#) in the *Amazon EC2 User Guide for Linux Instances*.

3. Connect to your Amazon EC2 Linux instance. For more information, see [Connect to your Linux instance](#) in the *Amazon EC2 User Guide for Linux Instances*.
4. Open a terminal on your EC2 instance using secure shell (SSH), and log in with the appropriate credentials.
5. Create a directory on the EC2 instance for mounting the SVM volume as follows:

```
sudo mkdir /fsx
```

6. Mount the volume to the directory you just created using the following command:

```
sudo mount -t nfs svm-dns-name:/volume-junction-path /fsx
```

The following example uses sample values.

```
sudo mount -t nfs svm-01234567890abcdef0.fs-01234567890abcdef1.fsx.us-east-1.amazonaws.com:/vol1 /fsx
```

You can also use the SVM's IP address SVM instead of its DNS name. We recommend using the DNS name to mount clients to scale-out file systems because it helps ensure that your clients are balanced across your file system's high-availability (HA) pairs.

```
sudo mount -t nfs 198.51.100.1:/vol1 /fsx
```

Note

For scale-out file systems, the parallel NFS (pNFS) protocol is enabled by default and is used by default for any clients mounting volumes with NFS v4.1 or greater.

Using `/etc/fstab` to mount automatically on instance reboot

To automatically remount your FSx for ONTAP volume when an Amazon EC2 Linux instance reboots, use the `/etc/fstab` file. The `/etc/fstab` file contains information about file systems. The command `mount -a`, which runs during instance start-up, mounts the file systems listed in `/etc/fstab`.

Note

FSx for ONTAP file systems do not support automatic mounting using `/etc/fstab` on Amazon EC2 Mac instances.

Note

Before you can update the `/etc/fstab` file of your EC2 instance, make sure that you already created your FSx for ONTAP file system. For more information, see [Creating FSx for ONTAP file systems](#).

To update the `/etc/fstab` file on your EC2 instance

1. Connect to your EC2 instance:
 - To connect to your instance from a computer running macOS or Linux, specify the `.pem` file for your SSH command. To do this, use the `-i` option and the path to your private key.
 - To connect to your instance from a computer running Windows, you can either use MindTerm or PuTTY. To use PuTTY, install it and convert the `.pem` file to a `.ppk` file.

For more information, see the following topics in the *Amazon EC2 User Guide for Linux Instances*:

- [Connecting to your Linux instance using SSH](#)
- [Connecting to your Linux instance from Windows using PuTTY](#)

2. Create a local directory that will be used to mount the SVM volume.

```
sudo mkdir /fsx
```

3. Open the `/etc/fstab` file in an editor of your choice.
4. Add the following line to the `/etc/fstab` file. Insert a tab character between each parameter. It should appear as one line with no line breaks.

```
svm-dns-name:volume-junction-path /fsx nfs nfsvers=version,defaults 0 0
```

You can also use the IP address of volume's SVM. The last three parameters indicate NFS options (which we set to default), dumping of file system and filesystem check (these are typically not used so we set them to 0).

5. Save the changes to the file.
6. Now mount the file share using the following command. The next time the system starts, the folder will be mounted automatically.

```
sudo mount /fsx  
sudo mount svm-dns-name:volume-junction-path
```

Your EC2 instance is now configured to mount the ONTAP volume whenever it restarts.

Mounting on Microsoft Windows clients

This section describes how to access data in your FSx for ONTAP file system with clients running the Microsoft Windows operating system. Review the following requirements, regardless of the type of client you are using.

This procedure assumes that the client and the file system are located in the same VPC and Amazon Web Services account. If the client is located on-premise or in a different VPC, Amazon Web Services account, or Amazon Web Services Region, this procedure also assumes that you've set up Amazon Transit Gateway or a dedicated network connection using Amazon Direct Connect or a private, secure tunnel using Amazon Virtual Private Network. For more information, see [Accessing data from outside the deployment VPC](#).

We recommend that you attach volumes to your Windows clients using the SMB protocol.

Prerequisites

To access an ONTAP storage volume using a Microsoft Windows client, you have to satisfy the following prerequisites:

- The SVM of the volume you are attaching must be joined to your organization's Active Directory, or you must be using a workgroup. For more information on joining your SVM to an Active Directory, see [Managing FSx for ONTAP storage virtual machines](#). For more information on using workgroups, see [Set up an SMB server in a workgroup overview](#) in the *NetApp Documentation Center*.
- The volume you are attaching has a security style setting of NTFS or mixed. For more information, see [Managing FSx for ONTAP volumes](#).

To attach an ONTAP volume on a Windows client using SMB and Active Directory

1. Open the Amazon EC2 console at <https://console.amazonaws.cn/ec2/>.
2. Create or select an Amazon EC2 instance running Microsoft Windows that is in the same VPC as the file system, and joined to the same Microsoft Active Directory as the volume's SVM.

For more information on launching an instance, see [Step 1: Launch an instance](#) in the *Amazon EC2 User Guide for Windows Instances*.

For more information about joining an SVM to an Active Directory, see [Managing FSx for ONTAP storage virtual machines](#).

3. Connect to your Amazon EC2 Windows instance. For more information, see [Connecting to your Windows instance](#) in the *Amazon EC2 User Guide for Windows Instances*.
4. Open a command prompt.
5. Run the following command. Replace the following:
 - Replace Z: with any available drive letter.
 - Replace DNS_NAME with the DNS name or the IP address of the SMB endpoint for the volume's SVM.
 - Replace SHARE_NAME with the name of an SMB share. C\$ is the default SMB share at the root of the SVM's namespace, but you shouldn't mount it as that exposes storage to the root volume and can cause security and service disruption. You should provide an SMB share name to mount instead of C\$. For more information about creating SMB shares, see [Managing SMB shares](#).

```
net use Z: \\DNS_NAME\SHARE_NAME
```


The following example uses sample values.

```
net use Z: \\corp.example.com\group_share
```

You can also use the IP address of the SVM instead of its DNS name. We recommend using the DNS name to mount clients to scale-out file systems because it helps ensure that your clients are balanced across your file system's high-availability (HA) pairs.

```
net use Z: \\198.51.100.5\group_share
```

Mounting on macOS clients

This section describes how to access data in your FSx for ONTAP file system with clients running the macOS operating system. Review the following requirements, regardless of the type of client you are using.

This procedure assumes that the client and the file system are located in the same VPC and Amazon Web Services account. If the client is located on-premise, or in a different VPC, Amazon Web Services account or Amazon Web Services Region, you've set up Amazon Transit Gateway or a dedicated network connection using Amazon Direct Connect or a private, secure tunnel using Amazon Virtual Private Network. For more information, see [Accessing data from outside the deployment VPC](#).

We recommend that you attach volumes to your Mac clients using the SMB protocol.

To mount an ONTAP volume on a macOS client using SMB

1. Open the Amazon EC2 console at <https://console.amazonaws.cn/ec2/>.
2. Create or select an Amazon EC2 Mac instance running the macOS that is in the same VPC as the file system.

For more information on launching an instance, see [Step 1: Launch an instance](#) in the *Amazon EC2 User Guide for Linux Instances*.

3. Connect to your Amazon EC2 Mac instance. For more information, see [Connect to your Linux instance](#) in the *Amazon EC2 User Guide for Linux Instances*.
4. Open a terminal on your EC2 instance using secure shell (SSH), and log in with the appropriate credentials.

5. Create a directory on the EC2 instance for mounting the volume as follows:

```
sudo mkdir /fsx
```

6. Mount the volume using the following command.

```
sudo mount -t smbfs filesystem-dns-name:/smb-share-name mount-point
```

The following example uses sample values.

```
sudo mount -t smbfs svm-01234567890abcde2.fs-01234567890abcde5.fsx.us-east-1.amazonaws.com:/C$ /fsx
```

You can also use the SVM's IP address instead of its DNS name. We recommend using the DNS name to mount clients to scale-out file systems because it helps ensure that your clients are balanced across your file system's high-availability (HA) pairs.

```
sudo mount -t smbfs 198.51.100.10:/C$ /fsx
```

C\$ is the default SMB share that you can mount to see the root of the SVM's namespace. If you've created any Server Message Block (SMB) shares in your SVM, provide the SMB share names instead of C\$. For more information about creating SMB shares, see [Managing SMB shares](#).

To mount an ONTAP volume on a macOS client using NFS

1. Open the Amazon EC2 console at <https://console.amazonaws.cn/ec2/>.
2. Create or select an Amazon EC2 instance running Amazon Linux 2 that is in the same VPC as the file system.

For more information on launching an EC2 Linux instance, see [Step 1: Launch an instance](#) in the *Amazon EC2 User Guide for Linux Instances*.

3. Connect to your Amazon EC2 Linux instance. For more information, see [Connect to your Linux instance](#) in the *Amazon EC2 User Guide for Linux Instances*.
4. Mount your FSx for ONTAP volume on the Linux EC2 instance by either using a user-data script during instance launch, or by running the following commands:

```
sudo mount -t nfs -o nfsvers=NFS_version svm-dns-name:/volume-junction-path /mount-point
```

The following example uses sample values.

```
sudo mount -t nfs -o nfsvers=4.1  
svm-01234567890abcdef0.fs-01234567890abcdef1.fsx.us-east-1.amazonaws.com:/vol1 /  
fsxontap
```

You can also use the SVM's IP address SVM instead of its DNS name. We recommend using the DNS name to mount clients to scale-out file systems because it helps ensure that your clients are balanced across your file system's HA pairs.

```
sudo mount -t nfs -o nfsvers=4.1 198.51.100.1:/vol1 /fsxontap
```

5. Mount the volume to the directory you just created using the following command.

```
sudo mount -t nfs svm-dns-name:/volume-junction-path /fsx
```

The following example uses sample values.

```
sudo mount -t nfs svm-01234567890abcdef0.fs-01234567890abcdef1.fsx.us-  
east-1.amazonaws.com:/vol1 /fsx
```

You can also use the SVM's IP address SVM instead of its DNS name. We recommend using the DNS name to mount clients to scale-out file systems because it helps ensure that your clients are balanced across your file system's high-availability (HA) pairs.

```
sudo mount -t nfs 198.51.100.1:/vol1 /fsx
```

Mounting iSCSI LUNs

Amazon FSx for NetApp ONTAP provides shared block storage support over the iSCSI (Internet Small Computer Systems Interface) protocol. You can enable iSCSI storage by provisioning LUNs (Logical Unit Number) and mapping them to initiator groups (igroups), exposing block storage to your Linux and Windows hosts.

Note

The iSCSI protocol isn't supported for FSx for ONTAP scale-out file systems, which are file systems with more than one high-availability (HA) pair of file servers.

Topics

- [Mounting iSCSI LUNs to a Linux client](#)
- [Mounting iSCSI LUNs to a Windows client](#)


Mounting iSCSI LUNs to a Linux client

The examples presented in these procedures use the following set up:

- The iSCSI LUN that is getting mounted to the Linux host is already created. For more information, see [Creating an iSCSI LUN](#).
- The Linux host that is mounting the iSCSI LUN is an Amazon EC2 instance running the Amazon Linux 2 Amazon Machine Image (AMI). It has VPC security groups configured to allow inbound and outbound traffic as described in [File System Access Control with Amazon VPC](#).
- The Linux host and the FSx for ONTAP file system are located in the same VPC and Amazon Web Services account. If the host is located in another VPC, you can use VPC peering or Amazon Transit Gateway to grant other VPCs access to the volume's iSCSI endpoints. For more information, see [Accessing data from outside the deployment VPC](#).

If you are using an EC2 instance running a different Linux AMI, some of the utilities that get installed on the host might be pre-installed, and you might use different commands to install required packages. Aside from installing packages, the commands used in this section are valid for other EC2 Linux AMIs.

We recommend that the EC2 instance be in the same availability zone as your file system's preferred subnet, as shown in the following graphic.



Topics

- [Install and configure iSCSI on the Linux client](#)
- [Configure iSCSI on the FSx for ONTAP file system](#)
- [Mount an iSCSI LUN on your Linux client](#)

Install and configure iSCSI on the Linux client

To install the iSCSI client

1. Confirm that `iscsi-initiator-utils` and `device-mapper-multipath` are installed on your Linux device. Connect to your Linux instance using an SSH client. For more information, see [Connect to your Linux instance using SSH](#).
2. Install `multipath` and the iSCSI client using the following command. Installing `multipath` is required if you want to automatically failover between your file servers.

```
~$ sudo yum install -y device-mapper-multipath iscsi-initiator-utils
```

3. To facilitate a faster response when automatically failing over between file servers when using multipath, set the replacement timeout value in the `/etc/iscsi/iscsid.conf` file to a value of 5 instead of using the default value of 120.

```
~$ sudo sed -i 's/node.session.timeo.replacement_timeout = .*/
node.session.timeo.replacement_timeout = 5/' /etc/iscsi/iscsid.conf; sudo cat /etc/
iscsi/iscsid.conf | grep node.session.timeo.replacement_timeout
```

4. Start the iSCSI service.

```
~$ sudo service iscsid start
```

Note that depending on your Linux version, you may have to use this command instead:

```
~$ sudo systemctl start iscsid
```

5. Confirm that the service is running using the following command.

```
~$ sudo systemctl status iscsid.service
```

The system responds with the following output:

```
iscsid.service - Open-iSCSI
   Loaded: loaded (/usr/lib/systemd/system/iscsid.service; disabled; vendor
   preset: disabled)
   Active: active (running) since Fri 2021-09-02 00:00:00 UTC; 1min ago
   Docs: man:iscsid(8)
        man:iscsiadm(8)
   Process: 14658 ExecStart=/usr/sbin/iscsid (code=exited, status=0/SUCCESS)
   Main PID: 14660 (iscsid)
   CGroup: /system.slice/iscsid.service
           ##14659 /usr/sbin/iscsid
           ##14660 /usr/sbin/iscsid
```

To configure iSCSI on your Linux client

1. To enable your clients to automatically failover between your file servers, you must configure multipath. Use the following command:

```
~$ sudo mpathconf --enable --with_multipathd y
```

2. Determine the initiator name of your Linux host using the following command. The location of the initiator name depends on your iSCSI utility. If you are using `iscsi-initiator-utils`, the initiator name is located in the file `/etc/iscsi/initiatorname.iscsi`.

```
~$ sudo cat /etc/iscsi/initiatorname.iscsi
```

The system responds with the initiator name.

```
InitiatorName=iqn.1994-05.com.redhat:abcdef12345
```

Configure iSCSI on the FSx for ONTAP file system

1. Connect to the NetApp ONTAP CLI on the FSx for ONTAP file system on which you created the iSCSI LUN using the following command. For more information, see [Using the NetApp ONTAP CLI](#).

```
~$ ssh fsxadmin@your_management_endpoint_ip
```

2. Create the initiator group (igroup) using the NetApp ONTAP CLI [lun igroup create](#) command. An initiator group maps to iSCSI LUNs and control which initiators (clients) have access to LUNs. Replace `host_initiator_name` with the initiator name from your Linux host that you retrieved in the previous procedure.

```
::> lun igroup create -vserver svm_name -igroup igroup_name -  
initiator host_initiator_name -protocol iscsi -ostype linux
```

If you want to make the LUNs mapped to this igroup available to multiple hosts, you can specify multiple initiator names separated with a comma. For more information, see [lun igroup create](#) in the *NetApp ONTAP Documentation Center*.

3. Confirm that the igroup exists using the [lun igroup show](#) command:

```
::> lun igroup show
```

The system responds with the following output:

Vserver	Igroup	Protocol	OS Type	Initiators
<i>svm_name</i>	<i>igroup_name</i>	iscsi	linux	iqn.1994-05.com.redhat:abcdef12345

4. This step assumes that you have already created an iSCSI LUN. If you have not, see [Creating an iSCSI LUN](#) for step-by-step instructions to do so.

Create a mapping from the LUN you created to the igroup you created, using the [lun mapping create](#), specifying the following attributes:

- *svm_name* – The name of the storage virtual machine providing the iSCSI target. The host uses this value to reach the LUN.
- *vol_name* – The name of the volume hosting the LUN.
- *lun_name* – The name that you assigned to the LUN.
- *igroup_name* – The name of the initiator group.
- *lun_id* – The LUN ID integer is specific to the mapping, not to the LUN itself. This is used by the initiators in the igroup as the Logical Unit Number use this value for the initiator when accessing the storage.

```
::> lun mapping create -vserver svm_name -path /vol/vol_name/lun_name -
igroup igroup_name -lun-id lun_id
```

5. Use the [lun show -path](#) command to confirm the LUN is created, online, and mapped.

```
::> lun show -path /vol/vol_name/lun_name -fields state,mapped,serial-hex
```

The system responds with the following output:

Vserver	Path	serial-hex	state	mapped
<i>svm_name</i>	/vol/ <i>vol_name</i> / <i>lun_name</i>	6c5742314e5d52766e796150	online	mapped

Save the `serial_hex` value (in this example, it is 6c5742314e5d52766e796150), you will use it in a later step to create a friendly name for the block device.

- Use the `network interface show -vserver` command to retrieve the addresses of the `iscsi_1` and `iscsi_2` interfaces for the SVM in which you've created your iSCSI LUN.

```
::> network interface show -vserver svm_name
```

The system responds with the following output:

Logical Current Is	Status	Network	Current
Vserver Port Home	Admin/Oper	Address/Mask	Node

<i>svm_name</i>			
iscsi_1	up/up	172.31.0.143/20	
FSxId0123456789abcdef8-01 e0e	true		
iscsi_2	up/up	172.31.21.81/20	
FSxId0123456789abcdef8-02 e0e	true		
nfs_smb_management_1	up/up	198.19.250.177/20	
FSxId0123456789abcdef8-01 e0e	true		

3 entries were displayed.

In this example, the IP address of `iscsi_1` is `172.31.0.143` and `iscsi_2` is `172.31.21.81`.

Mount an iSCSI LUN on your Linux client

- On your Linux client, use the following command to discover the target iSCSI nodes using `iscsi_1`'s IP address `iscsi_1_IP`.

```
~$ sudo iscsiadm --mode discovery --op update --type sendtargets --  
portal iscsi_1_IP
```

```
172.31.0.143:3260,1029  
iqn.1992-08.com.netapp:sn.9cfa2c41207a11ecac390182c38bc256:vs.3  
172.31.21.81:3260,1028  
iqn.1992-08.com.netapp:sn.9cfa2c41207a11ecac390182c38bc256:vs.3
```

In this example,

`iqn.1992-08.com.netapp:sn.9cfa2c41207a11ecac390182c38bc256:vs.3` corresponds to the `target_initiator` for the iSCSI LUN in the preferred availability zone.

2. (Optional) You can establish additional sessions with the `target_initiator`. Amazon EC2 has a bandwidth limit of 5 Gb/s (~625 MB/s) for single-flow traffic, but you can create multiple sessions to drive higher levels of throughput to your file system from a single client. For more information, see [Amazon EC2 instance network bandwidth](#) in the *Amazon Elastic Compute Cloud User Guide for Linux Instances*.

The following command establishes 8 sessions per initiator per ONTAP node in each availability zone, enabling the client to drive up to 40 Gb/s (5,000 MB/s) of aggregate throughput to the iSCSI LUN.

```
~$ sudo iscsiadm --mode node -T target_initiator --op update -n  
node.session.nr_sessions -v 8
```

3. Log into the target initiators. Your iSCSI LUNs are presented as available disks.

```
~$ sudo iscsiadm --mode node -T target_initiator --login
```

```
Logging in to [iface: default, target:  
iqn.1992-08.com.netapp:sn.9cfa2c41207a11ecac390182c38bc256:vs.3, portal:  
172.31.14.66,3260] (multiple)  
Login to [iface: default, target:  
iqn.1992-08.com.netapp:sn.9cfa2c41207a11ecac390182c38bc256:vs.3, portal:  
172.31.14.66,3260] successful.
```

The output above is truncated; you should see one `Logging in` and one `Login successful` response for each session on each file server. In the case of 4 sessions per node, there will be 8 `Logging in` and 8 `Login successful` responses.

4. Use the following command to verify that `dm-multipath` has identified and merged the iSCSI sessions by showing a single LUN with multiple policies. There should be an equal number of devices that are listed as active and those listed as enabled.

```
~$ sudo multipath -ll
```

In the output, the disk name is formatted as `dm-xyz`, where `xyz` is an integer. If there are no other multipath disks, this value is `dm-0`.

```
3600a09806c5742314e5d52766e79614f dm-xyz NETAPP ,LUN C-Mode
size=10G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle'
hwandler='0' wp=rw
|+- policy='service-time 0' prio=50 status=active
| |- 0:0:0:1 sda      8:0   active ready running
| |- 1:0:0:1 sdc      8:32  active ready running
| |- 3:0:0:1 sdg      8:96  active ready running
| ` - 4:0:0:1 sdh      8:112 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 2:0:0:1 sdb      8:16  active ready running
  |- 7:0:0:1 sdf      8:80  active ready running
  |- 6:0:0:1 sde      8:64  active ready running
  ` - 5:0:0:1 sdd      8:48  active ready running
```

Your block device is now connected to your Linux client. It is located under the path `/dev/dm-xyz`. You should not use this path for administrative purposes; instead, use the symbolic link that is under the path `/dev/mapper/wwid`, where `wwid` is a unique identifier for your LUN that is consistent across devices. In the next step, you'll provide a friendly name for the `wwid` so you can distinguish it from other multipathed disks.

To give your block device a friendly name

1. To provide your device a friendly name, create an alias in the `/etc/multipath.conf` file. To do this, add the following entry to the file using your preferred text editor, replacing the following placeholders:
 - Replace `serial_hex` with the value the you saved in the [Configure iSCSI on the FSx for ONTAP file system](#) procedure.
 - Add the prefix `3600a0980` to the `serial_hex` value as shown in the example. This is a unique preamble for the NetApp ONTAP distribution that Amazon FSx for NetApp ONTAP uses.
 - Replace `device_name` with the friendly name you want to use for your device.

```
multipaths {
```

```

multipath {
    wwid 3600a0980serial_hex
    alias device_name
}
}

```

As an alternative, you can copy and save the following script as a bash file, such as `multipath_alias.sh`. You can run the script with `sudo` privileges, replacing `serial_hex` (without the `3600a0980` prefix) and `device_name` with your respective serial number and the desired friendly name. This script searches for an uncommented `multipaths` section in the `/etc/multipath.conf` file. If one exists, it appends a `multipath` entry to that section; otherwise, it will create a new `multipaths` section with a `multipath` entry for your block device.

```

#!/bin/bash
SN=serial_hex
ALIAS=device_name
CONF=/etc/multipath.conf
grep -q '^multipaths {' $CONF
UNCOMMENTED=$?
if [ $UNCOMMENTED -eq 0 ]
then
    sed -i '/^multipaths {/a\tmultipath {\n\t\t\twwid 3600a0980'"${SN}"'\n\t\t\talias '"${ALIAS}"'\n\t\t}\n' $CONF
else
    printf "multipaths {\n\t\tmultipath {\n\t\t\t\t\twwid 3600a0980$SN\n\t\t\t\t\talias\n\t\t\t\t\t$ALIAS\n\t\t\t\t}\n}" >> $CONF
fi

```

- Restart the `multipathd` service for the changes to `/etc/multipathd.conf` take effect.

```
~$ systemctl restart multipathd.service
```

To partition the LUN

The next step is to format and partition your LUN using `fdisk`.

- Use the following command to verify that the path to your `device_name` is present.

```
~$ ls /dev/mapper/device_name
```

```
/dev/device_name
```

2. Partition the disk using `fdisk`. You'll enter an interactive prompt. Enter the options in the order shown. Note that the Last sector value will vary depending on the size of your iSCSI LUN (10GB in this example). You can make multiple partitions by using a value smaller than the last sector (20971519 in this example).

```
~$ sudo fdisk /dev/mapper/device_name
```

The `fdisk` interactive prompt starts.

```
Welcome to fdisk (util-linux 2.30.2).

Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.

Device does not contain a recognized partition table.
Created a new DOS disklabel with disk identifier 0x66595cb0.

Command (m for help): n
Partition type
   p primary (0 primary, 0 extended, 4 free)
   e extended (container for logical partitions)
Select (default p): p
Partition number (1-4, default 1): 1
First sector (2048-20971519, default 2048): 2048
Last sector, +sectors or +size{K,M,G,T,P} (2048-20971519, default
20971519): 20971519

Created a new partition 1 of type 'Linux' and of size 512 B.
Command (m for help): w
The partition table has been altered.
Calling ioctl() to re-read partition table.
Syncing disks.
```

After entering `w`, your new partition `/dev/mapper/partition_name` becomes available. The *partition_name* has the format `<device_name><partition_number>`. 1 was used as the partition number used in the `fdisk` command in the previous step.

3. Create your file system using `/dev/mapper/partition_name` as the path.

```
~$ sudo mkfs.ext4 /dev/mapper/partition_name
```

The system responds with the following output:

```
mke2fs 1.42.9 (28-Dec-2013)
Discarding device blocks: done
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=16 blocks
655360 inodes, 2621184 blocks
131059 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2151677952
80 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632
Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

To mount the LUN on the Linux client

1. Create a directory *directory_path* as the mount point for your file system.

```
~$ sudo mkdir /directory_path/mount_point
```

2. Mount the file system using the following command.

```
~$ sudo mount -t ext4 /dev/mapper/partition_name /directory_path/mount_point
```

3. (Optional) You can change the ownership of the mount directory to your user. Replace *username* with your username.

```
~$ sudo chown username:username /directory_path/mount_point
```

4. (Optional) Verify that you can read from and write data to the file system.

```
~$ echo "Hello world!" > /directory_path/mount_point/HelloWorld.txt
~$ cat directory_path/HelloWorld.txt
Hello world!
```

You have successfully created and mounted an iSCSI LUN on your Linux client.

Mounting iSCSI LUNs to a Windows client


The examples presented in these procedures use the following set up:

- The iSCSI LUN that is getting mounted to a Windows host is already created. For more information, see [Creating an iSCSI LUN](#).
- The Microsoft Windows host that is mounting the iSCSI LUN is an Amazon EC2 instance running a Microsoft Windows Server 2019 Amazon Machine Image (AMI). It has VPC security groups configured to allow inbound and outbound traffic as described in [File System Access Control with Amazon VPC](#).

You may be using a different Microsoft Windows AMI in your set up.

- The client and the file system are located in the same VPC and Amazon Web Services account. If the client is located in another VPC, you can use VPC peering or Amazon Transit Gateway to grant other VPCs access to the iSCSI endpoints. For more information, see [Accessing data from outside the deployment VPC](#).

We recommend that the EC2 instance be in the same availability zone as your file system's preferred subnet, as shown in the following graphic.



Topics

- [Configure iSCSI on the Windows client](#)
- [Configure iSCSI on the FSx for ONTAP file system](#)
- [Mount an iSCSI LUN on the Windows client](#)

Configure iSCSI on the Windows client

1. Use Windows Remote Desktop to connect to the Windows client on which you want to mount the iSCSI LUN. For more information, see [Connect to your Windows instance using RDP](#) in the *Amazon Elastic Compute Cloud User Guide*.
2. Open a Windows PowerShell as an Administrator. Use the following commands to enable iSCSI on your Windows instance and configure the iSCSI service to start automatically.

```
PS C:\> Start-Service MSiSCSI
PS C:\> Set-Service -Name msiscsi -StartupType Automatic
```


3. Retrieve the initiator name of your Windows instance. You'll use this value in configuring iSCSI on your FSx for ONTAP file system using the NetApp ONTAP CLI.

```
PS C:\> (Get-InitiatorPort).NodeAddress
```

The system responds with the initiator port:

```
iqn.1991-05.com.microsoft:ec2amaz-abc123d
```

4. To enable your clients to automatically failover between your file servers, you need install Multipath-I0 (MPIO) on your Windows instance. Use the following command:

```
PS C:\> Install-WindowsFeature Multipath-I0
```

5. Restart your Windows instance after the Multipath-I0 installation has completed. Keep your Windows instance open to perform steps for mounting the iSCSI LUN in a section that follows.

Configure iSCSI on the FSx for ONTAP file system

1. Connect to the NetApp ONTAP CLI on the FSx for ONTAP file system on which you created the iSCSI LUN using the following command. For more information, see [Using the NetApp ONTAP CLI](#).

```
~$ ssh fsxadmin@your_management_endpoint_ip
```

2. Using the NetApp ONTAP CLI [lun igroup create](#), create the initiator group, or igroup. An initiator group maps to iSCSI LUNs and controls which initiators (clients) have access to LUNs. Replace `host_initiator_name` with the initiator name from your Windows host that you retrieved in the previous procedure.

```
::> lun igroup create -vserver svm_name -igroup igroup_name -  
initiator host_initiator_name -protocol iscsi -ostype windows
```

If you want to make the LUNs mapped to this igroup available to multiple hosts, you can specify multiple comma-separated initiator names. For more information, see [lun igroup create](#) in the *NetApp ONTAP Documentation Center*.

3. Confirm the igroup was created successfully using the following command:

```
::> lun igroup show
```

The system responds with the following output:

Vserver	Igroup	Protocol	OS Type	Initiators
<i>svm_name</i>	<i>igroup_name</i>	iscsi	windows	iqn.1994-05.com.windows:abcdef12345

With the igroup created, you are ready to create LUNs and map them to the igroup.

- This step assumes that you have already created an iSCSI LUN. If you have not, see [Creating an iSCSI LUN](#) for step-by-step instructions to do so.

Create a LUN mapping from the LUN to your new igroup.

```
::> lun mapping create -vserver svm_name -path /vol/vol_name/lun_name -
igroup igroup_name -lun-id lun_id
```

- Confirm that the LUN is created, online, and mapped with the following command:

```
::> lun show -path /vol/vol_name/lun_name
```

Vserver	Path	State	Mapped	Type	Size
<i>svm_name</i>	<i>/vol/vol_name/lun_name</i>	online	mapped	windows	10GB

You are now ready to add the iSCSI target on your Windows instance.

- Retrieve the IP addresses of the `iscsi_1` and `iscsi_2` interfaces for your SVM using the following command:

```
::> network interface show -vserver svm_name
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
<i>svm_name</i>	<code>iscsi_1</code>	up/up	172.31.0.143/20	FSxId0123456789abcdef8-01	e0e	true
	<code>iscsi_2</code>	up/up	172.31.21.81/20	FSxId0123456789abcdef8-02	e0e	true

```

nfs_smb_management_1
      up/up      198.19.250.177/20  FSxId0123456789abcdef8-01
                                     e0e      true
3 entries were displayed.

```

In this example, the IP address of `iscsi_1` is `172.31.0.143` and `iscsi_2` is `172.31.21.81`.

Mount an iSCSI LUN on the Windows client

1. On your Windows instance, open a PowerShell terminal as an Administrator.
2. You will create a `.ps1` script that does the following:
 - Connects to each of your file system's iSCSI interfaces.
 - Adds and configures MPIO for iSCSI.
 - Establishes 8 sessions for each iSCSI connection, which enables the client to drive up to 40 Gb/s (5,000 MB/s) of aggregate throughput to the iSCSI LUN. Having 8 sessions ensures a single client can drive the full 4,000 MB/s throughput capacity for the highest-level FSx for ONTAP throughput capacity. You can optionally change the number of sessions to a higher or lower number of sessions (each session provides up to 625 MB/s of throughput) by modifying the script's for-loop in the `#Establish iSCSI connection` step from `1..8` to another upper-bound. For more information, see [Amazon EC2 instance network bandwidth](#) in the *Amazon Elastic Compute Cloud User Guide for Windows Instances*.

Copy the following set of commands into a file to create the `.ps1` script.

- Replace `iscsi_1` and `iscsi_2` with the IP addresses you retrieved in the previous step.
- Replace `ec2_ip` with the IP address of your Windows instance.

```

#iSCSI IP addresses for Preferred and Standby subnets
$TargetPortalAddresses = @("iscsi_1","iscsi_2")

#iSCSI Initiator IP Address (Local node IP address)
$LocaliSCSIAddress = "ec2_ip"

#Connect to FSx for NetApp ONTAP file system
Foreach ($TargetPortalAddress in $TargetPortalAddresses) {

```

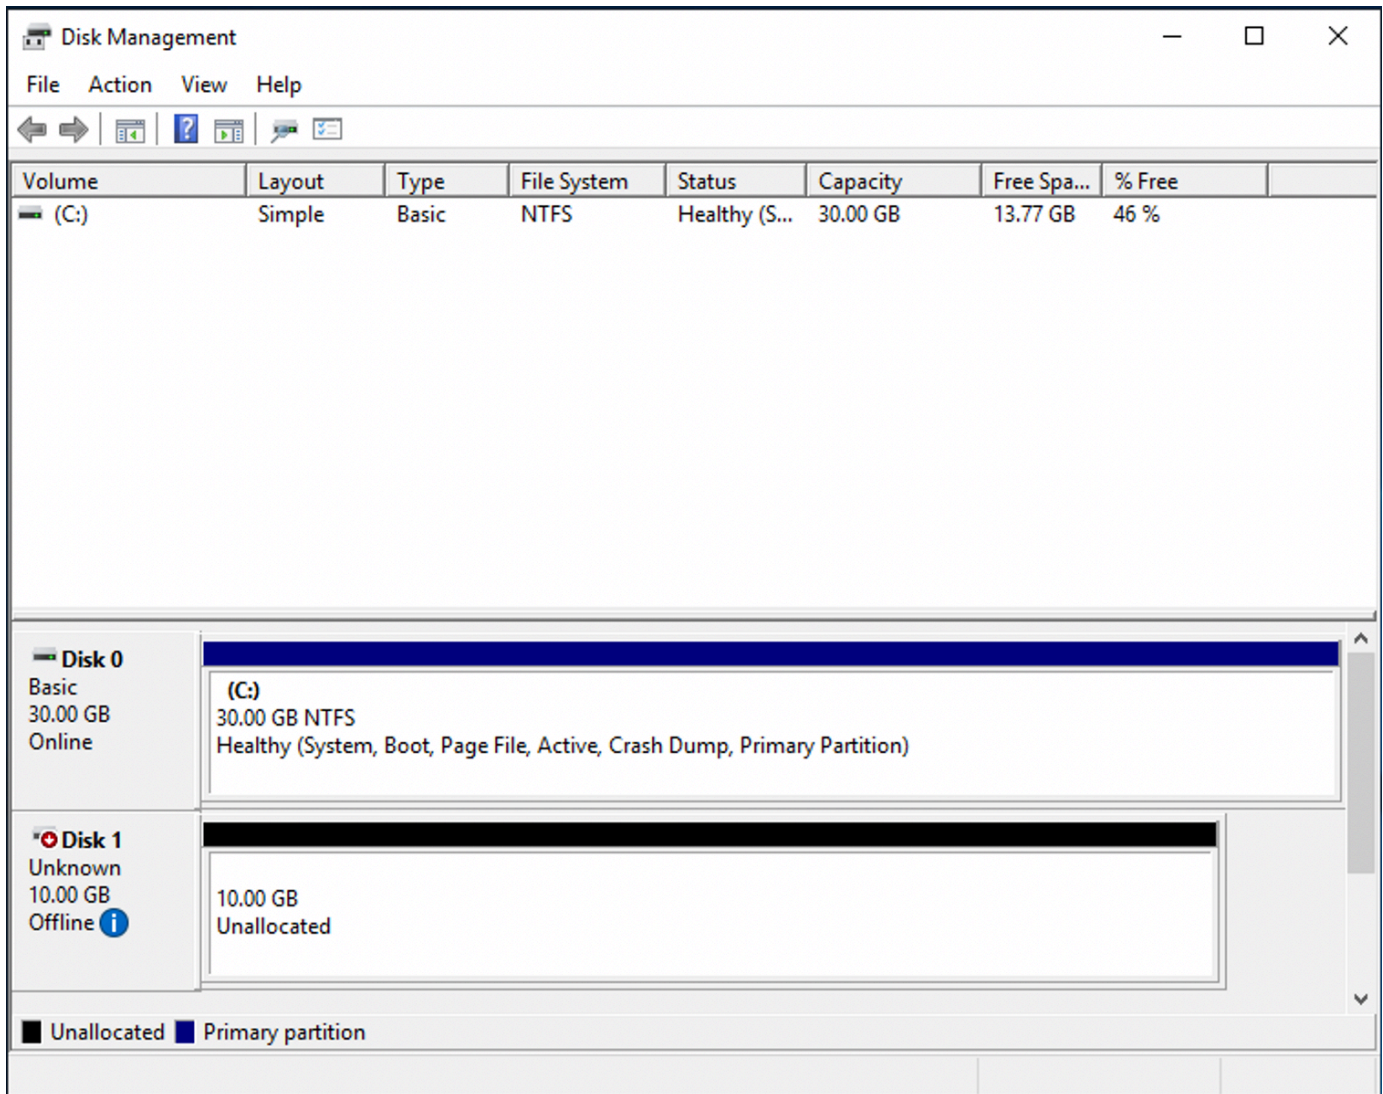
```
New-IscsiTargetPortal -TargetPortalAddress $TargetPortalAddress -
TargetPortalPortNumber 3260 -InitiatorPortalAddress $LocaliSCSIAddress
}

#Add MPIIO support for iSCSI
New-MSDSMSupportedHW -VendorId MSFT2005 -ProductId iSCSIBusType_0x9

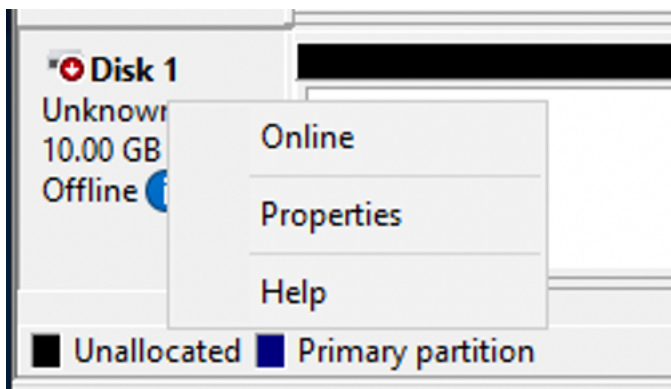
#Establish iSCSI connection
1..8 | %{Foreach($TargetPortalAddress in $TargetPortalAddresses)
{Get-IscsiTarget | Connect-IscsiTarget -IsMultipathEnabled $true -
TargetPortalAddress $TargetPortalAddress -InitiatorPortalAddress $LocaliSCSIAddress
-IsPersistent $true}}

#Set the MPIIO Policy to Round Robin
Set-MSDSMGlobalDefaultLoadBalancePolicy -Policy RR
```

3. Launch the Windows Disk Management application. Open the Windows Run dialog box, and enter `diskmgmt.msc` and press **Enter**. The Disk Management application opens.



4. Locate the unallocated disk This is the iSCSI LUN. In the example, Disk 1 is the iSCSI disk. It is offline.



Bring the volume online by placing the cursor over **Disk 1** and right-click then choose **Online**.

Note

You can modify the storage area network (SAN) policy so that new volumes are automatically brought online. For more information, see [SAN policies](#) in the *Microsoft Windows Server Command Reference*.

5. To initialize the disk, place the cursor over **Disk 1** right-click, and choose **Initialize**. The Initialize dialog appears. Choose **OK** initialize the disk.
6. Format the disk as you would normally. After formatting is complete, the iSCSI drive appears as a usable drive on the Windows client.

Using FSx for ONTAP with other Amazon services

In addition to Amazon EC2, you can use other Amazon services with your volumes to access your data.

Topics

- [Using Amazon WorkSpaces with FSx for ONTAP](#)
- [Using Amazon Elastic Container Service with FSx for ONTAP](#)
- [Using VMware Cloud with FSx for ONTAP](#)

Using Amazon WorkSpaces with FSx for ONTAP

FSx for ONTAP can be used with Amazon WorkSpaces to provide shared network-attached storage (NAS) or to store roaming profiles for Amazon WorkSpaces accounts. After connecting to an SMB file share with a WorkSpaces instance, the user can create and edit files on the file share.

The following procedures shows how to use Amazon FSx with Amazon WorkSpaces to provide roaming profile and home folder access a consistent experience and to provide a shared team folder for Windows and Linux Workspaces users. If you are new to Amazon WorkSpaces, you can create your first Amazon WorkSpaces environment with the instructions in [Get started with WorkSpaces Quick Setup](#) in the *Amazon WorkSpaces Administration Guide*.

Topics

- [Provide Roaming Profile support](#)

- [Provide a shared folder to access common files](#)

Provide Roaming Profile support

You can use Amazon FSx to provide Roaming Profile support to users in your organization. A user will have permissions to access only their Roaming Profile. The folder will be automatically connected using Active Directory Group Policies. With a Roaming Profile, users' data and desktop settings are saved when they log off an Amazon FSx file share enabling documents and settings to be shared between different WorkSpaces instances, and automatically backed up using Amazon FSx daily automatic backups.

Step 1: Create a profile folder location for domain users using Amazon FSx

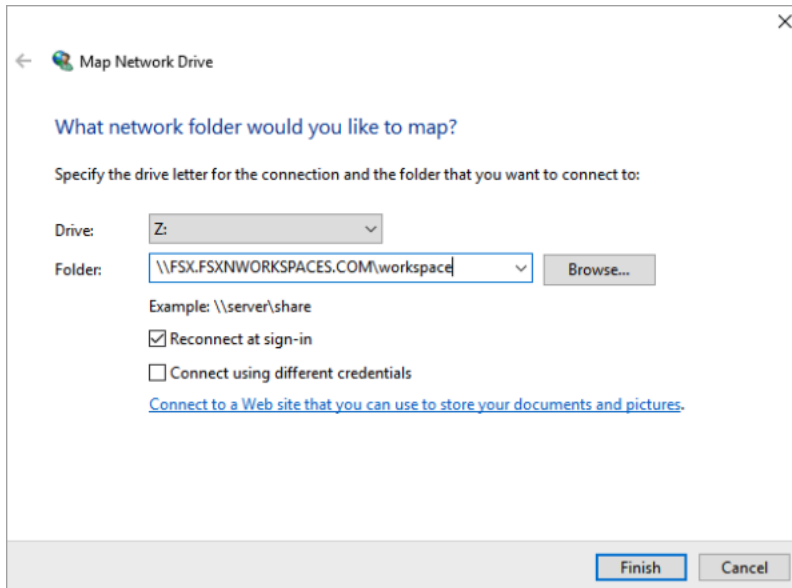
1. Create an FSx for ONTAP file system using the Amazon FSx console. For more information, see [To create a file system \(console\)](#).

Important

Each FSx for ONTAP file system has an endpoint IP address range from which the endpoints associated with the file system are created. For multi-AZ file systems, FSx for ONTAP chooses a default unused IP address range from 198.19.0.0/16 as the endpoint IP address range. This IP address range is also used by WorkSpaces for management traffic range, as described in [IP address and port requirements for WorkSpaces](#) in the *Amazon WorkSpaces Administration Guide*. As a result, to access your *multi-AZ* FSx for ONTAP file system from WorkSpaces, you must select an endpoint IP address range that does not overlap with 198.19.0.0/16.

2. If you don't have a storage virtual machine (SVM) joined to an Active Directory, create one now. For example, you can provision an SVM named `fsx` and set the security style to NTFS. For more information, see [To create a storage virtual machine \(console\)](#).
3. Create a volume for your SVM. For example, you can create a volume named `fsx-vol` which inherits the security style of your SVM's root volume. For more information, see [To create a FlexVol volume \(console\)](#).
4. Create an SMB share on your volume. For example, you can create a share called `workspace` on your volume named `fsx-vol`, in which you create a folder named `profiles`. For more information, see [Managing SMB shares](#).

5. Access your Amazon FSx SVM from an Amazon EC2 instance running Windows Server or from a WorkSpace. For more information, see [Accessing data](#).
6. You map your share to Z:\ on your Windows WorkSpaces instance:



Step 2: Link the FSx for ONTAP file share to User Accounts

1. On your test user's WorkSpace, choose **Windows > System > Advanced System Settings**.
2. In **System Properties**, select the **Advanced** tab and press the **Settings** button in the **User Profiles** section. The logged-in user will have a profile type of Local.
3. Log out the test user from the WorkSpace.
4. Set the test user to have a roaming profile located on your Amazon FSx file system. In your administrator WorkSpaces, open a PowerShell console and use a command similar to the following example (which uses the profiles folder you previously created in Step 1):


```
Set-ADUser username -ProfilePath \\filesystem-dns-name\sharename\foldername\username
```

For example:

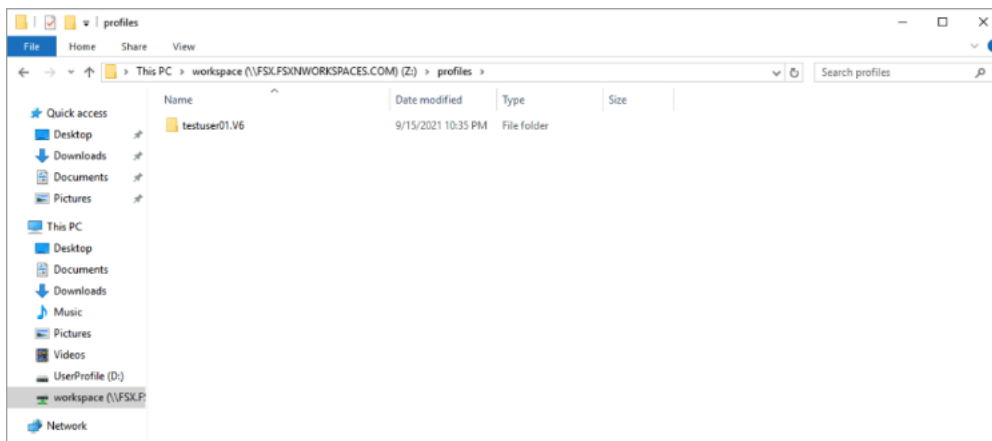
```
Set-ADUser testuser01 -ProfilePath \\fsx.fsxnworkspaces.com\workspace\profiles\testuser01
```

5. Log on to the test user WorkSpace.

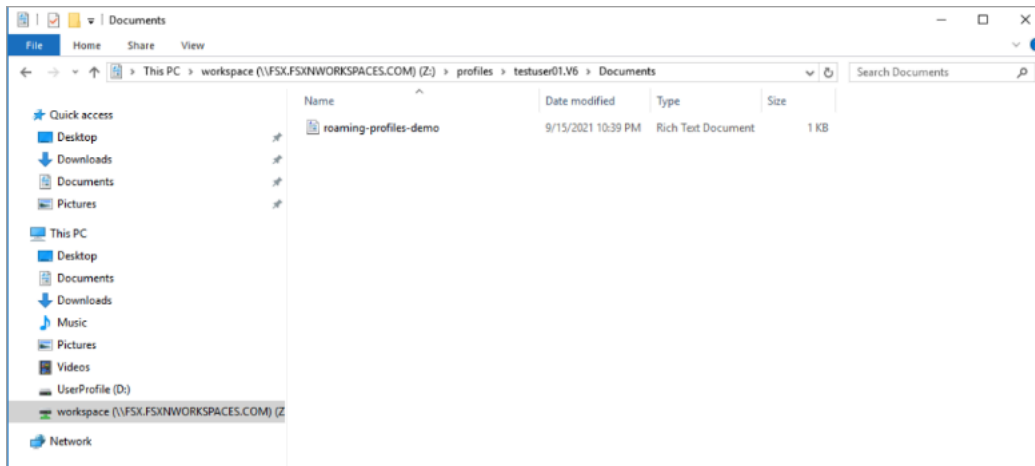
6. In **System Properties**, select the **Advanced** tab and press the **Settings** button in the **User Profiles** section. The logged-in user will have a profile type of **Roaming**.



7. Browse the FSx for ONTAP shared folder. In the **profiles** folder, you'll see a folder for the user.



8. Create a document in the test user's Documents folder
9. Log out the test user from their WorkSpace.
10. If you log back on as the test user and browse to their profile store, you will see the document you created.




Provide a shared folder to access common files


You can use Amazon FSx to provide a shared folder to users in your organization. A shared folder can be used to store files used by your user community, such as demo files, code examples, and instruction manuals needed by all users. Typically, you have drives mapped for shared folders; however because mapped drives use letters, there's a limit to the number of shares you can have. This procedure creates an Amazon FSx shared folder that's available without a drive letter, giving you greater flexibility in assigning shares to teams.

To mount a shared folder for cross-platform access from both Linux and Windows WorkSpaces

1. From the **Taskbar**, choose **Places > Connect to Server**.
 - a. For **Server**, enter *file-system-dns-name*.
 - b. Set **Type** to **Windows share**.
 - c. Set **Share** to the name of the SMB share, such as **workspace**.
 - d. You can leave **Folder** as **/** or set it to a folder, such as a folder named **team-shared**.
 - e. For a Linux WorkSpace, you don't need to enter your user details if your Linux WorkSpace is in the same domain as the Amazon FSx share.
 - f. Choose **Connect**.



2. After the connection is made, you can see the shared folder (named team-shared in this example) in the SMB share named workspace.



Using Amazon Elastic Container Service with FSx for ONTAP

You can access your Amazon FSx for NetApp ONTAP file systems from an Amazon Elastic Container Service (Amazon ECS) Docker container on an Amazon EC2 Linux or Windows instance.

Mounting on an Amazon ECS Linux container

1. Create an ECS cluster using the EC2 Linux + Networking cluster template for your Linux containers. For more information, see [Creating a cluster](#) in the *Amazon Elastic Container Service Developer Guide*.

2. Create a directory on the EC2 instance for mounting the SVM volume as follows:

```
sudo mkdir /fsxontap
```

3. Mount your FSx for ONTAP volume on the Linux EC2 instance by either using a user-data script during instance launch, or by running the following commands:

```
sudo mount -t nfs svm-ip-address:/vol1 /fsxontap
```

4. Mount the volume using the following command:

```
sudo mount -t nfs -o nfsvers=NFS_version svm-dns-name:/volume-junction-path /  
fsxontap
```

The following example uses sample values.

```
sudo mount -t nfs -o nfsvers=4.1  
svm-01234567890abcdef0.fs-01234567890abcdef1.fsx.us-east-1.amazonaws.com:/vol1 /  
fsxontap
```

You can also use the SVM's IP address SVM instead of its DNS name.

```
sudo mount -t nfs -o nfsvers=4.1 198.51.100.1:/vol1 /fsxontap
```

5. When creating your Amazon ECS task definitions, add the following volumes and mountPoints container properties in the JSON container definition. Replace the sourcePath with the mount point and directory in your FSx for ONTAP file system.

```
{
```

```
"volumes": [
  {
    "name": "ontap-volume",
    "host": {
      "sourcePath": "mountpoint"
    }
  }
],
"mountPoints": [
  {
    "containerPath": "containermountpoint",
    "sourceVolume": "ontap-volume"
  }
],
.
.
.
}
```

Mounting on an Amazon ECS Windows container

1. Create an ECS cluster using the EC2 Windows + Networking cluster template for your Windows containers. For more information, see [Creating a cluster](#) in the *Amazon Elastic Container Service Developer Guide*.
2. Add a domain-joined Windows EC2 instance to the ECS Windows cluster and map an SMB share.

Launch an ECS optimized Windows EC2 instance that is joined to your Active Directory domain and initialize the ECS agent by running the following command.

```
PS C:\Users\user> Initialize-ECSAgent -Cluster windows-fsx-cluster -
EnableTaskIAMRole
```

You can also pass the information in a script to the user-data text field as follows.

```
<powershell>
Initialize-ECSAgent -Cluster windows-fsx-cluster -EnableTaskIAMRole
</powershell>
```

3. Create an SMB global mapping on the EC2 instance so that you can map your SMB share to a drive. Replace the values below netbios or DNS name for your FSx file system and share name. The NFS volume vol1 that was mounted on the Linux EC2 instance is configured as a CIFS share fsxontap on the FSx file system.

```
vserver cifs share show -vserver svm08 -share-name fsxontap
```

```

                                Vserver: svm08
                                Share: fsxontap
CIFS Server NetBIOS Name: FSXONTAPDEMO
                                Path: /vol1
                                Share Properties: oplocks
                                                browsable
                                                changenotify
                                                show-previous-versions
                                Symlink Properties: symlinks
                                File Mode Creation Mask: -
                                Directory Mode Creation Mask: -
                                Share Comment: -
                                Share ACL: Everyone / Full Control
                                File Attribute Cache Lifetime: -
                                Volume Name: vol1
                                Offline Files: manual
                                Vscan File-Operations Profile: standard
                                Maximum Tree Connections on Share: 4294967295
                                UNIX Group for File Create: -

```

4. Create the SMB global mapping on the EC2 instance using the following command:

```
New-SmbGlobalMapping -RemotePath \\fsxontapdemo.fsxontap.com\fsxontap -LocalPath Z:
```

5. When creating your Amazon ECS task definitions, add the following volumes and mountPoints container properties in the JSON container definition. Replace the sourcePath with the mount point and directory in your FSx for ONTAP file system.

```
{
  "volumes": [
    {
      "name": "ontap-volume",
      "host": {
        "sourcePath": "mountpoint"
      }
    }
  ]
}
```

```
    }  
  }  
],  
"mountPoints": [  
  {  
    "containerPath": "containermountpoint",  
    "sourceVolume": "ontap-volume"  
  }  
],  
.  
.  
.  
}
```

Using VMware Cloud with FSx for ONTAP

You can use FSx for ONTAP as an external datastore for VMware Cloud on Amazon Software-Defined Data Centers (SDDCs). For more information, see [Configure Amazon FSx for NetApp ONTAP as External Storage](#) and [VMware Cloud on Amazon with Amazon FSx for NetApp ONTAP Deployment Guide](#).

Availability and durability

Amazon FSx for NetApp ONTAP uses two deployment types, Single-AZ and Multi-AZ, that offer different levels of availability and durability. This topic describes the availability and durability features of each deployment type to help you choose the one that is right for your workloads. For information on the service's availability SLA (Service Level Agreement), see [Amazon FSx Service Level Agreement](#).

Topics

- [Choosing a file system deployment type](#)
- [Failover process for FSx for ONTAP](#)
- [Network resources](#)

Note

Multi-AZ file systems are not available in the China (Beijing) region.

Choosing a file system deployment type

The availability and durability features of Single-AZ and Multi-AZ file system deployment types are described in the following sections.

Single-AZ deployment type

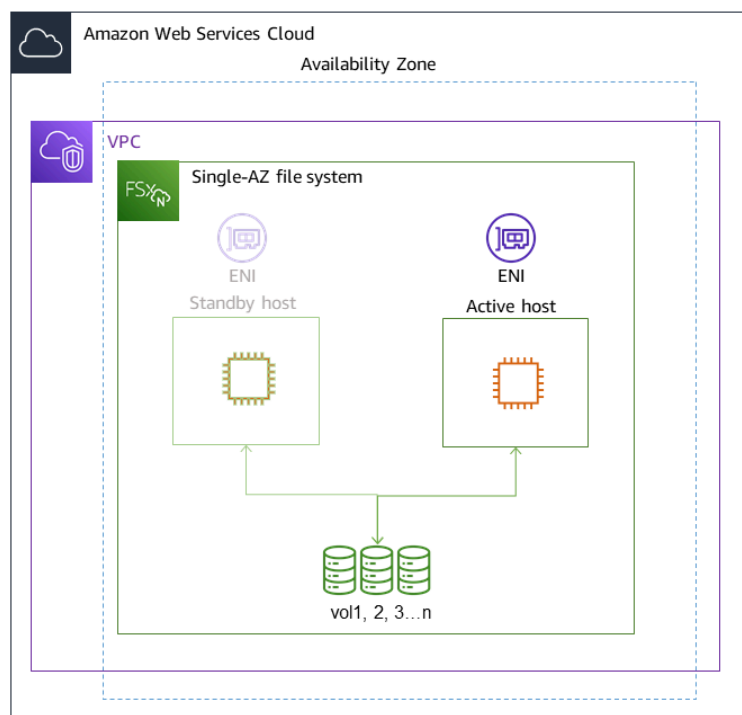
When you create a Single-AZ file system, Amazon FSx automatically provisions one to twelve pairs of file servers in an active-standby configuration, with the active and standby file servers in each pair located in separate fault domains within a single Availability Zone in the Amazon Web Services Region. During planned file system maintenance or an unplanned service disruption of any active file server, Amazon FSx automatically and independently fails over that high-availability (HA) pair to the standby file server, typically within a few seconds. During a failover, you continue to have access to your data without manual intervention.

To ensure high availability, Amazon FSx continuously monitors for hardware failures, and automatically replaces infrastructure components in the event of a failure. To achieve high

durability, Amazon FSx automatically replicates your data within an Availability Zone to protect it from component failure. In addition, you have the option to configure automatic daily backups of your file system data. These backups are stored across multiple Availability Zones to provide multi-AZ resiliency for all backup data.

Single-AZ file systems are designed for use cases that do not require the data resiliency model of a Multi-AZ file system. They provide a cost-optimized solution for use cases such as development and test environments, or storing secondary copies of data that is already stored on premises or in other Amazon Web Services Regions, by only replicating data within a single Availability Zone.

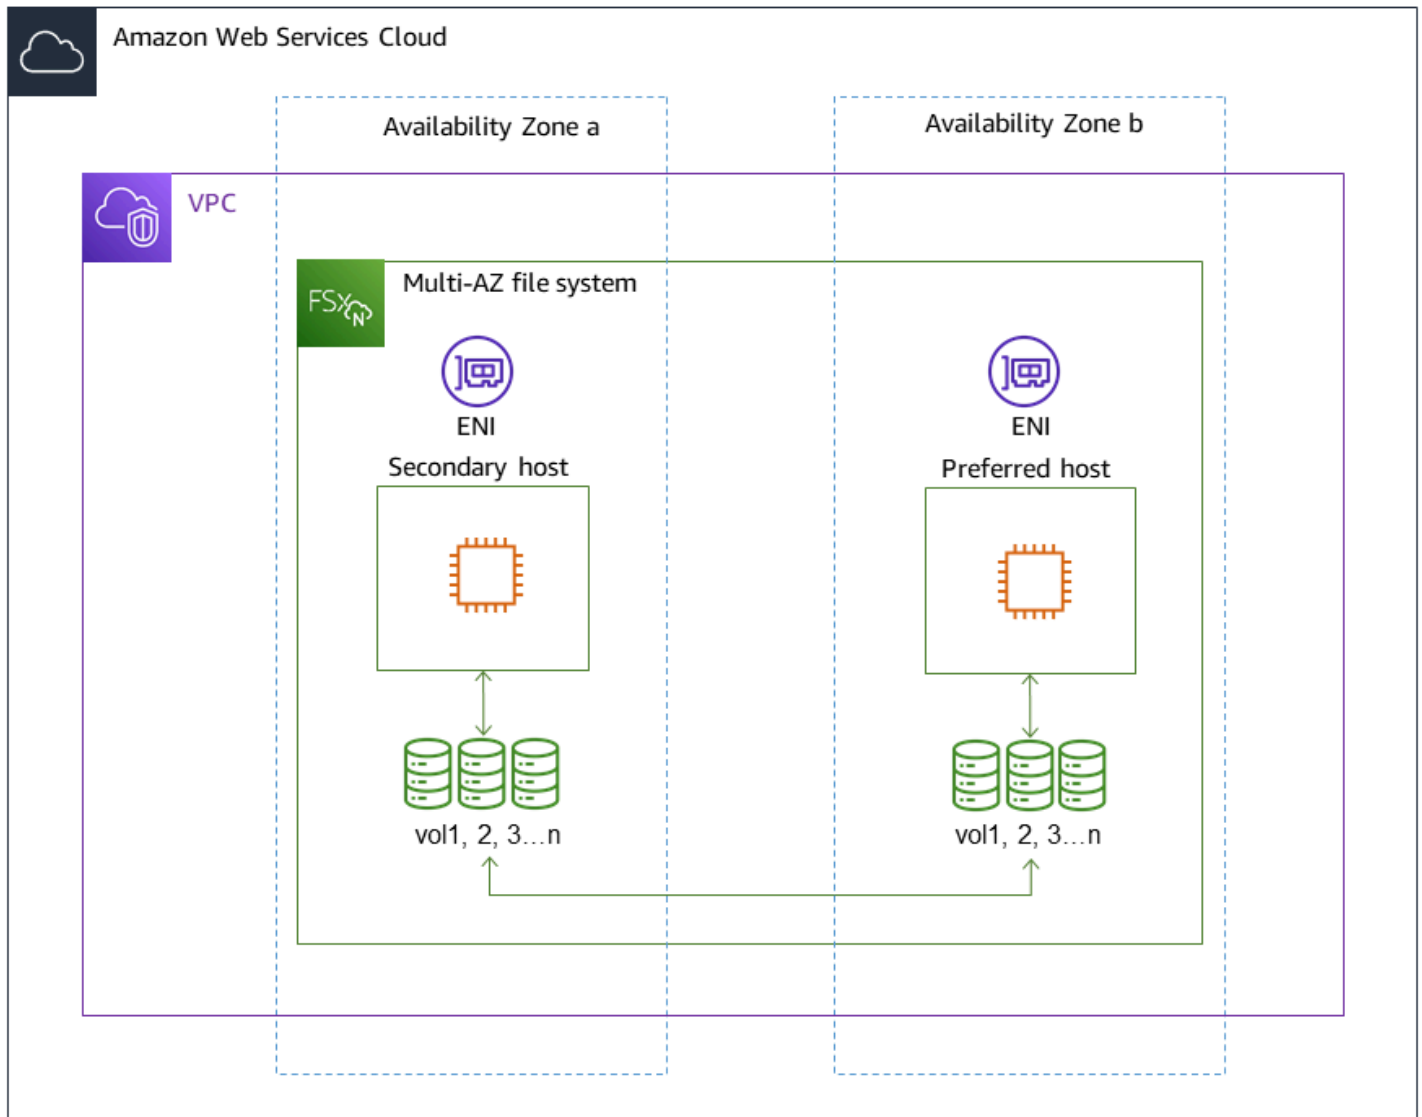
The following diagram illustrates the architecture for an FSx for ONTAP Single-AZ file system.



Multi-AZ deployment type

Multi-AZ file systems support all the availability and durability features of Single-AZ file systems. In addition, they are designed to provide continuous availability to data even when an Availability Zone is unavailable. Multi-AZ deployments have a single HA pair of file servers, the standby file server is deployed in a different Availability Zone from the active file server in the same Amazon Web Services Region. Any changes written to your file system are synchronously replicated across Availability Zones to the standby.

Multi-AZ file systems are designed for use cases such as business-critical production workloads that require high availability to shared ONTAP file data and need storage with built-in replication across Availability Zones. The following diagram illustrates the architecture for an FSx for ONTAP Multi-AZ file system.



Failover process for FSx for ONTAP

Single-AZ and Multi-AZ file systems automatically fail over a given HA pair from the preferred or active file server to the standby file server if any of the following conditions occur:

- The preferred or active file server becomes unavailable
- The file system's throughput capacity is changed

- The preferred or active file server undergoes planned maintenance
- An Availability Zone outage occurs (Multi-AZ file systems only)

Note

For scale-out file systems, each HA pair's failover behavior is independent. If the preferred file server for one HA pair is unavailable, only that HA pair will fail over to its standby file server.

When failing over from one file server to another, the new active file server automatically begins serving all file system read and write requests to that HA pair. For Multi-AZ file systems, when the preferred file server is fully recovered and becomes available, Amazon FSx automatically fails back to it, with failback usually completing in less than 60 seconds. For Single-AZ and Multi-AZ file systems, a failover typically completes in less than 60 seconds from the detection of the failure on the active file server to the promotion of the standby file server to active status. Because the endpoint IP address that clients use to access data over NFS or SMB remains the same, failovers are transparent to Linux, Windows, and macOS applications, which resume file system operations without manual intervention.

To ensure that failovers are transparent to clients connected to your FSx for ONTAP Single-AZ and Multi-AZ file systems, see [Accessing data from within Amazon](#).

Testing failover on a file system

You can test failover on your scale-up file system by modifying its throughput capacity. When you modify your file system's throughput capacity, Amazon FSx switches out the file system's file servers serially. File systems automatically fail over to the secondary server while Amazon FSx replaces the preferred file server first. Once updated, the file system automatically fails back to the new primary server and Amazon FSx replaces the secondary file server.

You can monitor the progress of the throughput capacity update request in the Amazon FSx console, the CLI, and the API. For more information about modifying your file system's throughput capacity and monitoring the progress of the request, see [Managing throughput capacity](#).

Network resources

This section describes the network resources consumed by Single-AZ and Multi-AZ file systems.

Subnets

When you create a Single-AZ file system, you specify a single subnet for the file system. The subnet you choose defines the Availability Zone in which the file system is created. When you create a Multi-AZ file system, you specify two subnets, one for the preferred file server, and one for the standby file server. The two subnets you choose must be in different Availability Zones within the same Amazon Web Services Region. For more information about Amazon VPC, see [What is Amazon VPC?](#) in the *Amazon Virtual Private Cloud User Guide*.

Note

Regardless of the subnet that you specify, you can access your file system from any subnet within the file system's VPC.

File system elastic network interfaces

For Single-AZ file systems, Amazon FSx provisions two [elastic network interfaces](#) (ENI) in the subnet that you associate with your file system. For Multi-AZ file systems, Amazon FSx also provisions two ENIs, one in each of the subnets that you associate with your file system. Clients communicate with your Amazon FSx file system using the elastic network interface. The network interfaces are considered to be within the service scope of Amazon FSx, despite being part of your account's VPC. Multi-AZ file systems use floating internet protocol (IP) addresses so that connected clients seamlessly transition between the preferred and standby file servers during a failover event.

Warning

- You must not modify or delete the elastic network interfaces associated with your file system. Modifying or deleting the network interface can cause a permanent loss of connection between your VPC and your file system.
- The elastic network interfaces associated with your file system will have routes automatically created and added to your default VPC and subnet route tables. Modifying or deleting these routes may cause temporary or permanent loss of connectivity for your file system clients.

The following table summarizes the subnet, elastic network interface, and IP address resources for each of the FSx for ONTAP file system deployment types:

	Single-AZ (scale-up)	Single-AZ (scale-out)	Multi-AZ (scale-up)
Number of subnets	1	1	2
Number of elastic network interfaces	2	2 per HA pair	2
Number of IP addresses per ENI	1 + the number of SVMs in the file system	HA pair count + HA pair count multiplied by the number of SVMs in the file system	1 + the number of SVMs in the file system
Number of VPC route table routes	N/A	N/A	1 + the number of SVMs in the file system

Once a file system or SVM is created, its IP addresses doesn't change until the file system is deleted.

⚠ Important

Amazon FSx doesn't support accessing file systems from, or exposing file systems to the public Internet. Amazon FSx automatically detaches any Elastic IP address which is a public IP address reachable from the Internet, that gets attached to a file system's elastic network interface.

Managing storage capacity

Amazon FSx for NetApp ONTAP provides a number of storage-related features you can use to manage storage capacity on your file system.

Topics

- [FSx for ONTAP storage tiers](#)
- [Choosing the right amount of file system SSD storage](#)
- [File system storage capacity and IOPS](#)
- [Volume storage capacity](#)

FSx for ONTAP storage tiers


Storage tiers are the physical storage media for an Amazon FSx for NetApp ONTAP file system. FSx for ONTAP offers the following storage tiers:

- *SSD tier* – The user-provisioned, high-performance solid-state drive (SSD) storage that's purpose-built for the active portion of your data set.
- *Capacity pool tier* – Fully elastic storage that automatically scales to petabytes in size, and is cost-optimized for your infrequently accessed data.

An FSx for ONTAP volume is a virtual resource that, similar to folders, doesn't consume storage capacity. The data that you store—and that consumes physical storage—lives inside volumes. When you create a volume, you specify its size—which you can modify after it's created. FSx for ONTAP volumes are thin provisioned, and file system storage is not reserved in advance. Instead, SSD and capacity pool storage are allocated dynamically, as needed. A [tiering policy](#), which you configure at the volume level, determines if and when data that's stored in the SSD tier transitions to the capacity pool tier.

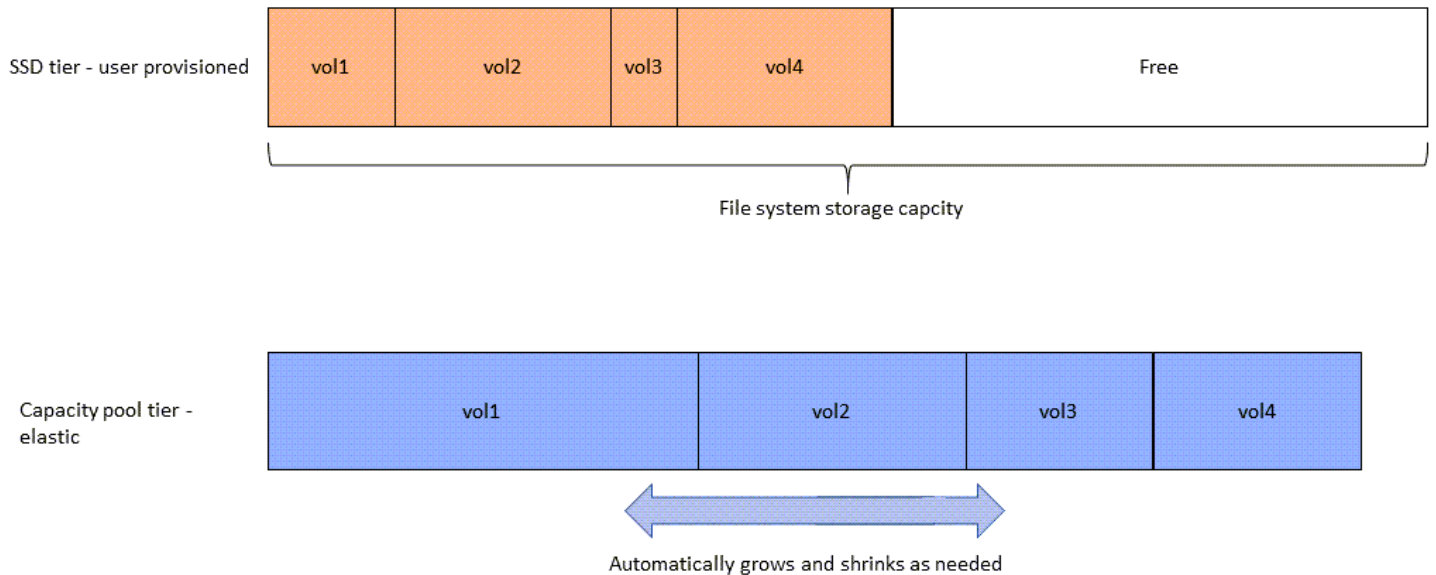
The following diagram illustrates an example of data laid out across multiple FSx for ONTAP volumes in a file system.

Volume thin provisioning



The following diagram illustrates how the file system's physical storage capacity is consumed by the data in the four volumes in the previous diagram.

Storage tiers – physical resource



You can reduce your storage costs by choosing the tiering policy that best meets the requirements for each volume on your file system. For more information, see [Volume data tiering](#).

Choosing the right amount of file system SSD storage

When choosing amount of SSD storage capacity for your FSx for ONTAP file system, you need to keep in mind the following items that impact the amount of SSD storage available for storing your data:

- Storage capacity reserved for the NetApp ONTAP software overhead.
- File metadata
- Recently written data
- Files that you intend to store on SSD storage, whether it's data that hasn't hit its cooling period, or data that you recently read that was retrieved back to SSD.

How SSD storage is used

Your file system's SSD storage is used for a combination of NetApp ONTAP software (overhead), file metadata, and your data.

NetApp ONTAP software overhead

Like other NetApp ONTAP file systems, up to 16% of a file system's SSD storage capacity is reserved for ONTAP overhead, which means it's not available for storing your files. The ONTAP overhead is allocated as follows:

- 11% is reserved for NetApp ONTAP software. For file systems with over 30 tebibytes (TiB) of SSD storage capacity, 6% is reserved.
- 5% is reserved for aggregate snapshots, which are required to synchronize data between both of a file system's file servers.

File metadata

File metadata typically consumes 3-7% of the storage capacity that is consumed by the files. This percentage depends on the average file size (a smaller average file size requires more metadata), and the amount of storage efficiency savings achieved on your files. Note that file metadata doesn't benefit from storage efficiency savings. You can use the following guidelines for estimating the amount of SSD storage used for metadata on your file system.

Average file size	Size of metadata as a percentage of file data
4 KB	7%
8 KB	3.5%
32 KB or greater	1-3%

When sizing the amount of SSD storage capacity you need for the metadata of files you plan to store on the capacity pool tier, we recommend using a conservative ratio of 1 GiB of SSD storage for every 10 GiB of data you plan to store on the capacity pool tier.

File data stored on your SSD tier

In addition to your active data set and all file metadata, all data written to your file system is initially written to the SSD tier before being tiered-off to capacity pool storage. This is true regardless of the volume's tiering policy, with the exception of transferring data using SnapMirror to a volume configured with an **All** data tiering policy.

Random reads from the capacity pool tier are cached in the SSD tier, as long as the SSD tier is under 90% utilization. For more information, see [Volume data tiering](#).

Recommended SSD capacity utilization

We recommend that you do not exceed 80% utilization of your SSD storage tier in an ongoing basis. For scale-out file systems, we additionally recommend that you don't exceed 80% utilization of any of your file system's aggregates on an ongoing basis. These recommendations is consistent with NetApp's recommendation for ONTAP. Because your file system's SSD tier is also used for staging writes to, and for random reads from, the capacity pool tier, any sudden changes in access patterns can quickly cause the utilization of your SSD tier to increase.

At 90% SSD utilization, data read from the capacity pool tier is no longer cached on the SSD tier so that the remaining SSD capacity is preserved for any new data that is written to the file system. This causes repeat reads of the same data from the capacity pool tier to be read from capacity pool storage instead of being cached and read from the SSD tier, which can impact the throughput capacity your file system.

All tiering functionality stops when the SSD tier is at or above 98% utilization. For more information, see [Tiering thresholds](#).

FSx for ONTAP storage efficiency

NetApp ONTAP offers block-level storage efficiency features including compression, compaction, and deduplication that can save you up to 65% in storage capacity for general file shares, without sacrificing performance.

Amazon FSx for NetApp ONTAP also supports other ONTAP features that save you space, including snapshots, thin provisioning, and FlexClone volumes.

Storage efficiency features are not enabled by default. You can enable them as follows:

- On an SVM's root volume when you [create a file system](#).
- When you [create a new volume](#).
- When you [modify an existing volume](#).

To view the amount of storage savings on a file system with storage efficiency enabled, see [Viewing storage efficiency savings](#).

Calculating storage efficiency savings

You can use the `LogicalDataStored` and `StorageUsed` FSx for ONTAP CloudWatch file system metrics to calculate storage savings from compression, deduplication, compaction, snapshots, and FlexClones. These metrics have a single dimension, `FileSystemId`. For more information, see [File system metrics](#).

- To compute storage-efficiency savings in bytes, take the Average of `StorageUsed` over a given period and subtract it from the Average of `LogicalDataStored` over the same period.
- To compute storage-efficiency savings as a percentage of total logical data size, take the Average of `StorageUsed` over a given period and subtract it from the Average of `LogicalDataStored` over the same period. Then divide the difference by the Average of `LogicalDataStored` over the same period.

SSD sizing example

Assume you want to store 100 TiB of data for an application where 80% of the data is infrequently accessed. In this scenario, 80% (80 TB) of your data is automatically tiered to the capacity pool tier and the remaining 20% (20 TB) remains in SSD storage. Based on the typical storage efficiency

savings of 65% for general-purpose file sharing workloads, that equates to 7 TiB of data. To maintain an 80% SSD utilization rate, you need 8.75 TiB of SSD storage capacity for the 20 TiB of actively-accessed data. The amount of SSD storage that you provision also needs to account for the ONTAP software storage overhead of 16%, as shown in the following calculation.

```
ssdNeeded = ssdProvisioned * (1 - 0.16)
8.75 TiB / 0.84 = ssdProvisioned
10.42 TiB = ssdProvisioned
```

So in this example, you need to provision at least 10.42 TiB of SSD storage. You will also use 28 TiB of capacity pool storage for the remaining 80 TiB of infrequently accessed data.

File system storage capacity and IOPS

When you create an FSx for ONTAP file system, you specify the storage capacity of the SSD tier. For scale-out file systems, the storage capacity that you specify is spread evenly among the storage pools of each high-availability (HA) pair; these storage pools are called *aggregates*.

For each GiB of SSD storage that you provision, Amazon FSx automatically provisions 3 SSD input/output operations per second (IOPS) for the file system, up to a maximum of 160,000 SSD IOPS per file system. For scale-out file systems, your SSD IOPS are spread evenly across each of your file system's aggregates. You have the option to specify a level of provisioned SSD IOPS above the automatic 3 SSD IOPS per GiB. For more information about the maximum number of SSD IOPS that you can provision for your FSx for ONTAP file system, see [Impact of throughput capacity on performance](#).

Topics

- [Updating file system SSD storage and IOPS](#)
- [Monitoring SSD storage utilization](#)
- [Setting a storage utilization alarm](#)
- [Viewing storage efficiency savings](#)
- [Modifying SSD storage capacity and provisioned IOPS](#)
- [Monitoring storage capacity and IOPS updates](#)
- [Increasing SSD storage capacity dynamically](#)

Updating file system SSD storage and IOPS

When you need additional storage for the active portion of your data set, you can increase the SSD storage capacity of your Amazon FSx for NetApp ONTAP file system. Use the Amazon FSx console, Amazon FSx API, or Amazon Command Line Interface (Amazon CLI) to increase the SSD storage capacity. For more information, see [Modifying SSD storage capacity and provisioned IOPS](#).

When you increase the SSD storage capacity of your Amazon FSx file system, the new capacity is typically available for use within minutes. You are billed for the new SSD storage capacity after it becomes available to you. For more information about pricing, see [Amazon FSx for NetApp ONTAP Pricing](#).

After you increase your storage capacity, Amazon FSx runs a storage-optimization process in the background to re-balance your data. For most file systems, storage optimization takes a few hours, with minimal noticeable impact to your workload performance.


You can track the progress of the storage-optimization process at any time by using the Amazon FSx console, CLI, and API. For more information, see [Monitoring storage capacity and IOPS updates](#).

Considerations

Here are a few important items to consider when modifying a file system's SSD storage capacity and provisioned IOPS:

- **Storage capacity increase only** – You can only *increase* the amount of SSD storage capacity for a file system; you can't decrease the storage capacity.
- **Storage capacity minimum increase** – Each SSD storage capacity increase must be a minimum of 10 percent of the file system's current SSD storage capacity, up to the maximum SSD storage capacity for your file system's configuration.
- **(Scale-out only) Storage capacity spread** – The new storage capacity or SSD IOPS that you select for your file system is spread evenly across each of your file system's aggregates.
- **Time between increases** – After modifying SSD storage capacity, provisioned IOPS, or throughput capacity on a file system, you must wait at least six hours before modifying any of these configurations on the same file system again. This is sometimes referred to as a cooldown period.
- **Provisioned IOPS modes** – For a provisioned IOPS change, you must specify one of the two IOPS modes:

- **Automatic** mode – Amazon FSx automatically scales your SSD IOPS to maintain 3 provisioned SSD IOPS per GiB of SSD storage capacity, up to the maximum SSD IOPS for your file system configuration.

 **Note**

For more information about the maximum number of SSD IOPS that you can provision for your FSx for ONTAP file system, see [Impact of throughput capacity on performance](#).

- **User-provisioned** mode – You specify the number of SSD IOPS, which must be greater than or equal to 3 IOPS per GiB of SSD storage capacity. If you choose to provision a higher level of IOPS, you pay for the average IOPS provisioned above your included rate for the month, measured in IOPS-months.

For more information about pricing, see [Amazon FSx for NetApp ONTAP Pricing](#).

When to increase SSD storage capacity

If you're running out of available SSD tier storage, we recommend that you increase the storage capacity of your file system. Running out of storage indicates that your SSD tier is undersized for the active portion of your data set.

To monitor the amount of free storage that's available on the file system, use the file system-level `StorageCapacity` and `StorageUsed` Amazon CloudWatch metrics. You can create a CloudWatch alarm on a metric and be notified when it drops below a specific threshold. For more information, see [Monitoring with Amazon CloudWatch](#).

 **Note**

We recommend that you don't exceed 80% SSD storage capacity utilization to ensure that data tiering, throughput scaling, and other maintenance activities function properly, and that there is capacity available for additional data. For scale-out file systems, this recommendation applies to both the average utilization across all of your file system's aggregates and to each individual aggregate.

For more information about how a file system's SSD storage is used and how much SSD storage is reserved for file metadata and operating software, see [Choosing the right amount of file system SSD storage](#).

Monitoring SSD storage utilization

You can monitor your file system's SSD storage capacity utilization using a variety of Amazon and NetApp tools. Using Amazon CloudWatch you can monitor storage capacity utilization and set alarms to alert you when storage capacity utilization reaches a customizable threshold.

Note

We recommend that you don't exceed 80% storage capacity utilization of your SSD storage tier. This ensures that tiering functions properly, and provides overhead for new data. If your SSD storage tier is consistently above 80% storage capacity utilization, you can increase your SSD storage tier's capacity. For more information, see [Updating file system SSD storage and IOPS](#).

You can view a file system's available SSD storage and the overall storage distribution in the Amazon FSx console. The **Available SSD storage capacity** graph displays the amount of available SSD-based storage capacity on a file system over time. The **Storage distribution** graph shows how a file system's overall storage capacity is currently distributed over 3 categories:

- Capacity pool tier
- SSD tier - available
- SSD tier - used

You can monitor your file system's SSD storage capacity utilization in the Amazon Web Services Management Console, using the following procedure.

To monitor file system available SSD tier storage capacity (console)

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. Choose **File systems** in the left-hand navigation column, then choose the ONTAP file system that you want to view storage capacity information for. The file system detail page appears.

- In the second panel, choose the **Monitoring & performance** tab, then choose **Storage**. The **Available primary storage capacity** and **Storage capacity utilization per aggregate** graphs are displayed.

Setting a storage utilization alarm

We recommend that you do not exceed an average SSD storage capacity utilization of 80% in an ongoing basis. Occasional SSD storage utilization spikes above 80% are acceptable. Maintaining an average utilization under 80% provides you with enough capacity to increase your storage without encountering issues. The following procedure shows how to create a CloudWatch alarm that alerts you to when your file system's SSD storage utilization is approaching 80%.

To create a primary storage utilization CloudWatch alarm

To calculate the average SSD storage-capacity utilization, use the `StorageUsed` metric. Divide by the `MaximumStorageCapacity` over the same period, with the `StorageTier` dimension equal to `SSD`.

- Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
- In the left navigation pane, choose **File systems**, and then choose the file system that you want to create the alarm for.
- On the **Summary** page, choose **Monitoring**.
- Choose **Create CloudWatch alarm**. You are redirected to the **Alarms > Create alarm > Specify metric and conditions** page in the CloudWatch console.
- Choose **Select metric**.
- In the **Metrics** section, choose **FSx**.
- Choose the **Detailed File System Metrics** metric category.
- To show only the metrics available for the file system that you're creating the alarm for, enter the file system's ID in the search field.
- For the file system that you want to set the alarm for, choose the following metrics:

Metric name	StorageTier	Data Type
StorageUsed	SSD	All
StorageCapacity	SSD	All

10. Select the **Graphed metrics** tab. For each of the metrics that you previously added, perform the following actions:
 - Set the value in the **Statistic** column for each the metrics to **Average**.
 - Set the **Period** value to one of the predefined evaluation periods.
11. Choose the **Add math** dropdown, and then select **Start with an empty expression** from the list of predefined metric math expressions.

After you choose **Start with an empty expression**, a math expression box appears where you apply or edit math expressions.

12. In the **Edit math expression** field, enter the expression for dividing the StorageUsed metric by the StorageCapacity metric, as follows:
 - Enter the label for the StorageUsed metric, for example m1.
 - Enter the forward slash character, /, for the division operation.
 - Enter the label for the StorageCapacity metric.

Choose **Apply**.

13. Clear the check boxes to the left of the metrics on the page. Only the check box next to the expression to use for the alarm should be selected. The expression that you choose for the alarm must produce a single time series and show only one line on the graph. Then choose **Select metric**.

The **Specify metric and conditions** page appears, showing a graph and other information about the math expression that you have selected.


14. For **Whenever *expression* is**, specify that the expression must be greater than, less than, or equal to the threshold. Under **than...**, specify the threshold value.

To create an alarm that notifies you when the SSD storage capacity is approaching the 80% threshold, set **Whenever *expression* is** equal to the threshold, and specify a threshold value of 80%.

15. Choose **Additional configuration**. For **Datapoints to alarm**, specify how many evaluation periods (data points) must be in the ALARM state to trigger the alarm. If the two values here match, you create an alarm that goes to ALARM state if that many consecutive periods are breaching.

To create an M out of N alarm, specify a lower number for the first value than you specify for the second value. For more information, see [Evaluating an alarm](#) in the *Amazon CloudWatch User Guide*.


16. For **Missing data treatment**, choose how to have the alarm behave when some data points are missing. To prevent unnecessary and misleading alarm condition changes and to configure your alarms so that they are resilient to missing data points, see [Configuring how CloudWatch alarms treat missing data](#) in the *Amazon CloudWatch User Guide*.

 **Note**

Metrics may not be published during file system maintenance.

17. Choose **Next**.
18. Under **Notification**, select an SNS topic to notify when the alarm is in ALARM state, OK state, or INSUFFICIENT_DATA state.

If you choose **Create topic**, you can set the name and email addresses for a new email subscription list. This list is saved and appears in the field for future alarms.

 **Note**

If you use **Create topic** to create a new Amazon SNS topic, the email addresses must be verified before they receive notifications. Emails are sent only when the alarm enters an alarm state. If this alarm state change happens before the email addresses are verified, they do not receive a notification.

To have the alarm send multiple notifications for the same alarm state or for different alarm states, choose **Add notification**.

To have the alarm not send notifications, choose **Remove**.

If you want CloudWatch to send you an email or Amazon SNS notification when the alarm state initiates the action, choose an alarm state for **Whenever this alarm state is**.

19. When finished, choose **Next**.
20. Enter a name and description for the alarm. The name must contain only ASCII characters. Then choose **Next**.

21. On the **Preview and create** page, review the alarm that you're about to create, and then choose **Create Alarm**.

Viewing storage efficiency savings

When enabled, you can see how much storage capacity you are saving in the Amazon FSx console, the Amazon CloudWatch console, and the ONTAP CLI.

To view storage efficiency savings (console)

The storage efficiency savings displayed in the Amazon FSx console for an FSx for ONTAP file system includes the savings from FlexClones and Snapshots.

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. Choose the FSx for ONTAP file system that you want to view storage efficiency saving for from the list of **File systems**.
3. Choose **Summary** in the **Monitoring & performance** tab on the second panel in the file system details page.
4. The **Storage efficiency savings** chart shows how much space you are saving as a percentage of your logical data size and in physical bytes.

To view storage efficiency savings (ONTAP CLI)

You can see storage efficiency savings of just compaction, compression, and deduplication – without the effects of snapshots and FlexClones – by running the `storage aggregate show-efficiency` command using the ONTAP CLI. For more information, see [storage aggregate show-efficiency](#) in the NetApp ONTAP Documentation Center.

1. To access the NetApp ONTAP CLI, establish an SSH session on the management port of the Amazon FSx for NetApp ONTAP file system by running the following command. Replace *management_endpoint_ip* with the IP address of the file system's management port.

```
[~]$ ssh fsxadmin@management_endpoint_ip
```

For more information, see [Managing file systems with the ONTAP CLI](#).

2. The **storage aggregate show-efficiency** command displays information about the storage efficiency of all the aggregates. The storage efficiency is displayed at four different levels:

- Total
- Aggregate
- Volume
- Snapshot and FlexClone volume

```
::*> aggr show-efficiency
```

```
Aggregate: aggr1
Node: node1
```

```
Total Data Reduction Efficiency Ratio: 3.29:1
```

```
Total Storage Efficiency Ratio: 4.29:1
```

```
Aggregate: aggr2
```

```
Node: node1
```

```
Total Data Reduction Efficiency Ratio: 4.50:1
```

```
Total Storage Efficiency Ratio: 5.49:1
```

```
cluster::*> aggr show-efficiency -details
```

```
Aggregate: aggr1
Node: node1
```

```
Total Data Reduction Ratio: 2.39:1
```

```
Total Storage Efficiency Ratio: 4.29:1
```

```
Aggregate level Storage Efficiency
```

```
(Aggregate Deduplication and Data Compaction): 1.00:1
```

```
Volume Deduplication Efficiency: 5.03:1
```

```
Compression Efficiency: 1.00:1
```

```
Snapshot Volume Storage Efficiency: 8.81:1
```

```
FlexClone Volume Storage Efficiency: 1.00:1
```

```
Number of Efficiency Disabled Volumes: 1
```

```
Aggregate: aggr2
Node: node1
```

```
Total Data Reduction Ratio: 2.39:1
```

```
Total Storage Efficiency Ratio: 4.29:1
```

```
Aggregate level Storage Efficiency
(Aggregate Deduplication and Data Compaction): 1.00:1
Volume Deduplication Efficiency:                5.03:1
Compression Efficiency:                        1.00:1

Snapshot Volume Storage Efficiency:             8.81:1
FlexClone Volume Storage Efficiency:           1.00:1
Number of Efficiency Disabled Volumes:         1
```

Modifying SSD storage capacity and provisioned IOPS

You can increase a file system's SSD-based storage, and you increase or decrease the amount of provisioned SSD IOPS by using the Amazon FSx console, the Amazon CLI, and the API.

To update SSD storage capacity or provisioned IOPS for a file system (console)

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. In the left navigation pane, choose **File systems**. In the **File systems** list, select the FSx for ONTAP file system that you want to update SSD storage capacity and SSD IOPS for.
3. Choose **Actions > Update storage capacity**. Or, in the **Summary** section, choose **Update** next to the file system's **SSD storage capacity** value.

The **Update SSD storage capacity and IOPS** dialog box appears.

Update SSD storage capacity and IOPS ✕

File system ID

fs-01234567890abcdef

Current configuration

SSD storage capacity: 4096 GiB

IOPS mode: Automatic (3 IOPS per GiB of SSD storage)

SSD IOPS: 12288

SSD storage capacity

Modify storage capacity

Input type

Percentage

Absolute

Desired % increase

%

Minimum 4506 GiB (10% above current); Maximum 1048576 GiB.

Provisioned SSD IOPS

Automatic (3 IOPS per GiB of SSD storage)

User-provisioned

Configuration preview

Attribute	Current configuration	New configuration
SSD storage capacity	4,096 GiB (2,048 GiB per HA pair)	4,506 GiB (2,253 GiB per HA pair)
	Mode: Automatic	Mode: Automatic

4. To increase SSD storage capacity, choose **Modify storage capacity**.
5. For **Input type**, choose one of the following:
 - To enter the new SSD storage capacity as a percentage change from the current value, choose **Percentage**.
 - To enter the new value in GiB, choose **Absolute**.
6. Depending on the input type, enter a value for **Desired % increase**.
 - For **Percentage**, enter the percentage increase value. This value must be at least 10 percent greater than the current value.
 - For **Absolute**, enter the new value in GiB, up to the maximum allowed value of 196,608 GiB.
7. For **Provisioned SSD IOPS**, you have two options to modify the number of provisioned SSD IOPS for your file system:
 - If you want Amazon FSx to automatically scale your SSD IOPS to maintain 3 provisioned SSD IOPS per GiB of SSD storage capacity (up to a maximum of 160,000), choose **Automatic**.
 - If you want to specify the number of SSD IOPS, choose **User-provisioned**. Enter an absolute number of IOPS that's at least three times the amount of GiB of your SSD storage tier, and less than or equal to 160,000.

Note

For more information about the maximum number of SSD IOPS that you can provision for your FSx for ONTAP file system, see [Impact of throughput capacity on performance](#).

8. Choose **Update**.

Note

At the bottom of the prompt, a configuration preview is shown for your new SSD storage capacity and SSD IOPS. For scale-out file systems, the per-HA-pair value is also shown.

To update SSD storage capacity and provisioned IOPS for a file system (CLI)

To update the SSD storage capacity and provisioned IOPS for an FSx for ONTAP file system, use the Amazon CLI command [update-file-system](#) or the equivalent [UpdateFileSystem](#) API action. Set the following parameters with your values:

- Set `--file-system-id` to the ID of the file system that you are updating.
- To increase your SSD storage capacity, set `--storage-capacity` to the target storage capacity value, which must be at least 10 percent greater than the current value.
- To modify your provisioned SSD IOPS, use the `--ontap-configuration` `DiskIopsConfiguration` property. This property has two parameters, `Iops` and `Mode`:
 - If you want to specify the number of provisioned IOPS, use `Iops=number_of_IOPS` (up to a maximum of 160,000) and `Mode=USER_PROVISIONED`. The IOPS value must be greater than or equal to three times the requested SSD storage capacity. If you're not increasing the storage capacity, the IOPS value must be greater than or equal to three times the current SSD storage capacity.
 - If you want Amazon FSx to automatically increase your SSD IOPS, use `Mode=AUTOMATIC` and don't use the `Iops` parameter. Amazon FSx will automatically maintain 3 SSD IOPS per GiB of the provisioned SSD storage capacity (up to a maximum of 160,000).

Note

For more information about the maximum number of SSD IOPS that you can provision for your FSx for ONTAP file system, see [Impact of throughput capacity on performance](#).

The following example increases the file system's SSD storage to 2000 GiB and sets amount of user provisioned SSD IOPS to 7000.

```
aws fsx update-file-system \  
--file-system-id fs-0123456789abcdef0 \  
--storage-capacity 2000 \  
--ontap-configuration 'DiskIopsConfiguration={Iops=7000,Mode=USER_PROVISIONED}'
```

To monitor the progress of the update, use the [describe-file-systems](#) Amazon CLI command. Look for the `AdministrativeActions` section in the output.


For more information, see [AdministrativeAction](#) in the *Amazon FSx for NetApp ONTAP API Reference*.

Monitoring storage capacity and IOPS updates

You can monitor the progress of an SSD storage capacity and IOPS update by using the Amazon FSx console, CLI, and API.

To monitor storage and IOPS updates (console)

In the **Updates** tab on the **File system details** page for your FSx for ONTAP file system, you can view the 10 most recent updates for each update type.



The screenshot shows the 'Updates (2)' section in the Amazon FSx console. It features a search bar labeled 'Filter updates', a pagination control showing '1' of 1 items, and a settings icon. Below is a table with the following data:

Update type	Target value	Status	Progress %	Request time
Throughput capacity	256	Completed	-	2022-03-12T12:16:46-05:00
Storage capacity	1127	Updated; Optimizing	-	2022-03-12T12:17:02-05:00

For SSD storage capacity and IOPS updates, you can view the following information:

Update type

Supported types are **Storage capacity**, **Mode**, and **IOPS**. The **Mode** and **IOPS** values are listed for all storage capacity and IOPS scaling requests.

Target value

The value that you specified to update the file system's SSD storage capacity or IOPS to.

Status

The current status of the update. The possible values are as follows:

- **Pending** – Amazon FSx received the update request, but hasn't started processing it.
- **In progress** – Amazon FSx is processing the update request.
- **Updated; Optimizing** – Amazon FSx increased the file system's SSD storage capacity. The storage-optimization process is now rebalancing your data in the background.
- **Completed** – The update finished successfully.

- **Failed** – The update request failed. Choose the question mark (?) to see details.

Progress %

Displays the progress of the storage-optimization process as the percentage complete.

Request time

The time that Amazon FSx received the update action request.

To monitor storage and IOPS updates (CLI)

You can view and monitor file system SSD storage capacity increase requests by using the [describe-file-systems](#) Amazon CLI command and the [DescribeFileSystems](#) API operation. The `AdministrativeActions` array lists the 10 most recent update actions for each administrative action type. When you increase a file system's SSD storage capacity, two `AdministrativeActions` actions are generated: a `FILE_SYSTEM_UPDATE` and a `STORAGE_OPTIMIZATION` action.

The following example shows an excerpt of the response of a `describe-file-systems` CLI command. The file system has a pending administrative action to increase the SSD storage capacity to 2000 GiB and the provisioned SSD IOPS to 7000.

```
"AdministrativeActions": [  
  {  
    "AdministrativeActionType": "FILE_SYSTEM_UPDATE",  
    "RequestTime": 1586797629.095,  
    "Status": "PENDING",  
    "TargetFileSystemValues": {  
      "StorageCapacity": 2000,  
      "OntapConfiguration": {  
        "DiskIopsConfiguration": {  
          "Mode": "USER_PROVISIONED",  
          "Iops": 7000  
        }  
      }  
    }  
  },  
  {  
    "AdministrativeActionType": "STORAGE_OPTIMIZATION",  
    "RequestTime": 1586797629.095,  
    "Status": "PENDING"
```

```
  }  
]
```

Amazon FSx processes the `FILE_SYSTEM_UPDATE` action first, adding the new larger storage disks to the file system. When the new storage is available to the file system, the `FILE_SYSTEM_UPDATE` status changes to `UPDATED_OPTIMIZING`. The storage capacity shows the new larger value, and Amazon FSx begins processing the `STORAGE_OPTIMIZATION` administrative action. This behavior is shown in the following excerpt of the response of a `describe-file-systems` CLI command.

The `ProgressPercent` property displays the progress of the storage-optimization process. After the storage-optimization process has completed successfully, the status of the `FILE_SYSTEM_UPDATE` action changes to `COMPLETED`, and the `STORAGE_OPTIMIZATION` action no longer appears.

```
"AdministrativeActions": [  
  {  
    "AdministrativeActionType": "FILE_SYSTEM_UPDATE",  
    "RequestTime": 1586799169.445,  
    "Status": "UPDATED_OPTIMIZING",  
    "TargetFileSystemValues": {  
      "StorageCapacity": 2000,  
      "OntapConfiguration": {  
        "DiskIopsConfiguration": {  
          "Mode": "USER_PROVISIONED",  
          "Iops": 7000  
        }  
      }  
    }  
  },  
  {  
    "AdministrativeActionType": "STORAGE_OPTIMIZATION",  
    "ProgressPercent": 41,  
    "RequestTime": 1586799169.445,  
    "Status": "IN_PROGRESS"  
  }  
]
```

If the storage capacity or IOPS update request fails, the status of the `FILE_SYSTEM_UPDATE` action changes to `FAILED`, as shown in the following example. The `FailureDetails` property provides information about the failure.

```
"AdministrativeActions": [  
  {  
    "AdministrativeActionType": "FILE_SYSTEM_UPDATE",  
    "RequestTime": 1586373915.697,  
    "Status": "FAILED",  
    "TargetFileSystemValues": {  
      "StorageCapacity": 2000,  
      "OntapConfiguration": {  
        "DiskIopsConfiguration": {  
          "Mode": "USER_PROVISIONED",  
          "Iops": 7000  
        }  
      }  
    },  
    "FailureDetails": {  
      "Message": "failure-message"  
    }  
  }  
]
```

Increasing SSD storage capacity dynamically

You can use the following solution to dynamically increase the SSD storage capacity of an FSx for ONTAP file system when the amount of used SSD storage capacity exceeds a threshold that you specify. This Amazon CloudFormation template automatically deploys all of the components that are required to define the storage capacity threshold, the Amazon CloudWatch alarm based on this threshold, and the Amazon Lambda function that increases the file system's storage capacity.

The solution automatically deploys all of the components needed, and uses the following parameters:

- Your FSx for ONTAP file system ID.
- The used SSD storage capacity threshold (numerical value). This is the percentage at which the CloudWatch alarm will be triggered.
- The percentage by which to increase the storage capacity (%).
- The email address used to receive scaling notifications.


Topics

- [Architecture overview](#)

- [Amazon CloudFormation template](#)
- [Automated deployment with Amazon CloudFormation](#)

Architecture overview

Deploying this solution builds the following resources in the Amazon Web Services Cloud.



The diagram illustrates the following steps:

1. The Amazon CloudFormation template deploys a CloudWatch alarm, an Amazon Lambda function, an Amazon Simple Notification Service (Amazon SNS) queue, and all required Amazon Identity and Access Management (IAM) roles. The IAM role gives the Lambda function permission to invoke the Amazon FSx API operations.
2. CloudWatch triggers an alarm when the file system's used storage capacity exceeds the specified threshold, and sends a message to the Amazon SNS queue. An alarm is triggered only when the file system's used capacity exceeds the threshold continuously for a 5-minute period.
3. The solution then triggers the Lambda function that is subscribed to this Amazon SNS topic.
4. The Lambda function calculates the new file system storage capacity based on the specified percent increase value and sets the new file system storage capacity.

5. The original CloudWatch alarm state and results of the Lambda function operations are sent to the Amazon SNS queue.

To receive notifications about the actions that are performed as a response to the CloudWatch alarm, you must confirm the Amazon SNS topic subscription by following the link provided in the **Subscription Confirmation** email.

Amazon CloudFormation template

This solution uses Amazon CloudFormation to automate deploying the components that are used to automatically increase the storage capacity of an FSx for ONTAP file system. To use this solution, download the [FSxOntapDynamicStorageScaling](#) Amazon CloudFormation template.

The template uses the **Parameters** described as follows. Review the template parameters and their default values, and modify them for the needs of your file system.

FileSystemId

No default value. The ID of the file system for which you want to automatically increase the storage capacity.

LowFreeDataStorageCapacityThreshold

No default value. Specifies the used storage capacity threshold at which to trigger an alarm and automatically increase the file system's storage capacity, specified in percentage (%) of the file system's current storage capacity. The file system is considered to have low free storage capacity when the used storage exceeds this threshold.

EmailAddress

No default value. Specifies the email address to use for the SNS subscription and receives the storage capacity threshold alerts.

PercentIncrease

Default is **20%**. Specifies the amount by which to increase the storage capacity, expressed as a percentage of the current storage capacity.

Note

Storage scaling is attempted once every time the CloudWatch alarm enters the ALARM state. If your SSD storage capacity utilization remains above the threshold after a

storage scaling operation is attempted, the storage scaling operation isn't attempted again.

MaxFSxSizeinGiB

Default is **196608**. Specifies the maximum supported storage capacity for the SSD storage.

Automated deployment with Amazon CloudFormation

The following procedure configures and deploys an Amazon CloudFormation stack to automatically increase the storage capacity of an FSx for ONTAP file system. It takes a few minutes to deploy. For more information about creating a CloudFormation stack, see [Creating a stack on the Amazon CloudFormation console](#) in the *Amazon CloudFormation User Guide*.

Note

Implementing this solution incurs billing for the associated Amazon services. For more information, see the pricing details pages for those services.

Before you start, you must have the ID of the Amazon FSx file system that's running in the Amazon Virtual Private Cloud (Amazon VPC) in your Amazon Web Services account. For more information about creating Amazon FSx resources, see [Getting started with Amazon FSx for NetApp ONTAP](#).

To launch the automatic storage capacity increase solution stack

1. Download the [FSxOntapDynamicStorageScaling](#) Amazon CloudFormation template.

Note

Amazon FSx is currently only available in specific Amazon Regions. You must launch this solution in an Amazon Region where Amazon FSx is available. For more information, see [Amazon FSx endpoints and quotas](#) in the *Amazon Web Services General Reference*.

2. From the Amazon CloudFormation console, choose **Create stack > With new resources**.
3. Choose **Template is ready**. In the **Specify template** section, choose **Upload a template file** and upload the template that you downloaded.

4. In **Specify stack details**, enter the values for your automatic storage capacity increase solution.

Stack name

Stack name

Stack name can include letters (A-Z and a-z), numbers (0-9), and dashes (-).

Parameters

Parameters are defined in your template and allow you to input custom values when you create or update a stack.

Dynamic Storage Scaling Parameters

File system ID
Amazon FSx file system ID

Threshold
Used storage capacity threshold (%)

Percentage Capacity Increase
The percentage increase in storage capacity when used storage exceeds LowFreeDataStorageCapacityThreshold. Minimum increase is 10 %

Email address
The email address for alarm notification.

Maximum supported file system storage capacity (DO NOT MODIFY)
Maximum size supported for the primary SSD storage tier.

Cancel Previous Next

5. Enter a **Stack name**.
6. For **Parameters**, review the parameters for the template and modify them to meet the needs of your file system. Then choose **Next**.

 **Note**

To receive email notifications when scaling is attempted by this CloudFormation template, confirm the SNS subscription email that you receive after deploying the template.

7. Enter the **Options** settings that you want for your custom solution, and then choose **Next**.
8. For **Review**, review and confirm the solution settings. You must select the check box acknowledging that the template creates IAM resources.
9. Choose **Create** to deploy the stack.

You can view the status of the stack in the Amazon CloudFormation console in the **Status** column. You should see a status of **CREATE_COMPLETE** in a few minutes.

Updating the stack

After the stack is created, you can update it by using the same template and providing new values for the parameters. For more information, see [Updating stacks directly](#) in the *Amazon CloudFormation User Guide*.

Volume storage capacity

FSx for ONTAP volumes are virtual resources that you use for grouping data, determining how data is stored, and determining the type of access to your data. Volumes, like folders, don't consume file system storage capacity themselves. Only the data that's stored in a volume consumes SSD storage and, depending on the [volume's tiering policy](#), capacity pool storage. You set a volume's size when you create it, and you can change its size later. You can monitor and manage the storage capacity of your FSx for ONTAP volumes using the Amazon Web Services Management Console, Amazon CLI and API, and the ONTAP CLI.

Topics

- [Volume data tiering](#)
- [Snapshots and volume storage capacity](#)
- [Volume file capacity](#)
- [Updating a volume's storage capacity](#)
- [Enabling volume autosizing](#)
- [Monitoring volume storage capacity](#)
- [Setting a volume's tiering policy](#)
- [Setting the minimum cooling days](#)
- [Setting a volume's cloud retrieval policy](#)
- [Viewing a volume's file capacity](#)
- [Increasing the maximum number of files on a volume](#)
- [Enabling a volume's cloud write mode](#)

Volume data tiering

An Amazon FSx for NetApp ONTAP file system has two storage tiers: primary storage and capacity pool storage. Primary storage is provisioned, scalable, high-performance SSD storage that's purpose-built for the active portion of your data set. Capacity pool storage is a fully elastic storage tier that can scale to petabytes in size and is cost-optimized for infrequently accessed data.

The data on each volume is automatically tiered to the capacity pool storage tier based on the volume's tiering policy, cooling period, and threshold settings. The following sections describe ONTAP volume tiering policies and the thresholds used to determine when data is tiered to the capacity pool.

Volume tiering policies

You determine how to use your FSx for ONTAP file system's storage tiers by choosing the tiering policy for each of volume on the file system. You choose the tiering policy when you create a volume, and you can modify it at any time with the Amazon FSx console, Amazon CLI, API, or using [NetApp management tools](#). You can choose from one of the following policies that determine which data, if any, is tiered to the capacity pool storage.

Note

Tiering can move your file data and snapshot data to the capacity pool tier. However, file metadata always remains on the SSD tier. For more information, see [How SSD storage is used](#).

- **Auto** – This policy moves all cold data—user data and snapshots—to the capacity pool tier. The cooling rate of data is determined by the policy's cooling period, which by default is 31 days, and is configurable to values between 2–183 days. When the underlying cold data blocks are read randomly (as in typical file access), they are made hot and written to the primary storage tier. When cold data blocks are read sequentially (for example, by an antivirus scan), they remain cold and remain on the capacity pool storage tier. This is the default policy when creating a volume using the Amazon FSx console.
- **Snapshot Only** – This policy moves only snapshot data to the capacity pool storage tier. The rate at which snapshots are tiered to the capacity pool is determined by the policy's cooling period, which by default is set to 2 days, and is configurable to values between 2–183 days. When cold snapshot data are read, they are made hot and written to the primary storage tier. This is the

default policy when creating a volume using the Amazon CLI, Amazon FSx API, or the NetApp ONTAP CLI.

- **All** – This policy marks all user data and snapshot data as cold, and stores it in the capacity pool tier. When data blocks are read, they remain cold and are not written to the primary storage tier. When data is written to a volume with the **All** tiering policy, it is still initially written to the SSD storage tier, and is tiered to the capacity pool by a background process. Note that file metadata always remains on the SSD tier.
- **None** – This policy keeps all of your volume's data on the primary storage tier, and prevents it from being moved to capacity pool storage. If you set a volume to this policy after it used any other policy, existing data in the volume that was in capacity pool storage is moved to SSD storage by a background process as long as your SSD utilization is below 90%. This background process can be sped up by intentionally reading data or by modifying your volume's cloud retrieval policy. For more information, see [Cloud retrieval policies](#).

As a best practice, when migrating data that you plan to store long-term in capacity pool storage, we recommend that you use the **Auto** tiering policy on your volume. With **Auto** tiering, data is stored on the SSD storage tier for a minimum of 2 days (based on the volume's cooling period) before it's moved to the capacity pool tier. Retaining data on SSD storage for at least 2 days allows ONTAP to run post-process compression and deduplication savings on your data, which are preserved when data is tiered to the capacity pool. ONTAP only runs post-process compression and deduplication for data on SSD storage, so selecting this policy can help you maximize your long-term storage savings. You can also maximize the transfer speeds of the first backups you create of your volumes, as the data that is being backed up is on SSD storage.

For more information about setting or modifying a volume's tiering policy, see [Setting a volume's tiering policy](#).

Tiering cooling period

A volume's tiering cooling period sets the amount of time that it takes for data in the SSD tier to be marked as cold. The cooling period applies to the Auto and Snapshot-only tiering policies. You can set the cooling period to a value in the range of 2–183 days. For more information about setting the cooling period, see [Setting the minimum cooling days](#).

Data is tiered 24–48 hours after its cooling period expires. Tiering is a background process that consumes network resources, and has a lower priority than client-facing requests. Tiering activities are throttled when there are ongoing client-facing requests.

Cloud retrieval policies

A volume's cloud retrieval policy sets the conditions that specify when data that's read from the capacity pool tier is allowed to be promoted to the SSD tier. When the cloud retrieval policy is set to anything other than `Default`, this policy overrides the retrieval behavior of your volume's tiering policy. A volume can have one of the following cloud retrieval policies:

- **Default** – This policy retrieves tiered data based on the volume's underlying tiering policy. This is the default cloud retrieval policy for all volumes.
- **Never** – This policy never retrieves tiered data, regardless of whether the reads are sequential or random. This is similar to setting the tiering policy of your volume to **All**, except that you can use it with other policies—**Auto**, **Snapshot-only**—to tier data according to the minimum cooling period instead of immediately.
- **On-read** – This policy retrieves tiered data for all client-driven data reads. This policy has no effect when using the **All** tiering policy.
- **Promote** – This policy marks all of a volume's data that's in the capacity pool for retrieval to the SSD tier. The data is marked the next time the daily background tiering scanner runs. This policy is beneficial for applications that have cyclical workloads that run infrequently, but require SSD tier performance when they do run. This policy has no effect when using the **All** tiering policy.

For information on setting a volume's cloud retrieval policy, see [Setting a volume's cloud retrieval policy](#).

Tiering thresholds

A file system's SSD storage capacity utilization determines how ONTAP manages the tiering behavior for all of your volumes. Based on a file system's SSD storage capacity usage, the following thresholds set the tiering behavior as described. For information about how to monitor the capacity utilization of a volume's SSD storage tier, see [Monitoring volume storage capacity](#).

Note

We recommend that you don't exceed 80% storage capacity utilization of your SSD storage tier. For scale-out file systems, this recommendation applies to both the total average utilization across all of your file system's aggregates and to the utilization of each individual aggregate. This ensures that tiering functions properly, and provides overhead for new data. If your SSD storage tier is consistently above 80% storage capacity utilization,

you can increase your SSD storage tier's capacity. For more information, see [Updating file system SSD storage and IOPS](#).

FSx for ONTAP uses the following storage capacity thresholds to manage tiering on volumes:

- **<=50% SSD storage tier utilization** – At this threshold, the SSD storage tier is considered to be underutilized, and only volumes that are using the **All** tiering policy have data tiered to capacity pool storage. Volumes with **Auto** and **Snapshot-only** policies don't tier data at this threshold.
- **> 50% SSD storage tier utilization** – Volumes with **Auto** and **Snapshot-only** tiering policies tier data based on the tiering minimum cooling days setting. The default setting is 31 days.
- **>=90% SSD storage tier utilization** – At this threshold, Amazon FSx prioritizes preserving space in the SSD storage tier. Cold data from the capacity pool tier is no longer moved into the SSD storage tier when read for volumes using **Auto** and **Snapshot-only** policies.
- **>=98% SSD storage tier utilization** – All tiering functionality stops when the SSD storage tier is at or above 98% utilization. You can continue to read from storage tiers, but you can't write to the tiers.

Snapshots and volume storage capacity

A *snapshot* is a read-only image of an Amazon FSx for NetApp ONTAP volume at a point in time. Snapshots offer protection against accidental deletion or modification of files in your volumes. With snapshots, your users can easily view and restore individual files or folders from an earlier snapshot.

Snapshots are stored alongside your file system's data, and they consume the file system's storage capacity. However, snapshots consume storage capacity only for the portions of files that changed since the last snapshot. Snapshots are not included in backups of your file system volumes.

Snapshots are enabled by default on your volumes, using the default snapshot policy. Snapshots are stored in the `.snapshot` directory at the root of a volume. You can manage volume storage capacity for snapshots in the following ways:

- [Snapshot policies](#) – Select a built-in snapshot policy or choose a custom policy that you created in the ONTAP CLI or REST API.
- [Manually delete snapshots](#) – Reclaim storage capacity by deleting snapshots manually.

- [Create a snapshot autodelete policy](#) – Create a policy that deletes more snapshots than the default snapshot policy.
- [Turn off automatic snapshots](#) – Conserve storage capacity by turning off automatic snapshots.

For more information, see [Working with snapshots](#).

Volume file capacity

Amazon FSx for NetApp ONTAP volumes have file pointers that are used to store file metadata such as file name, last accessed time, permissions, size, and to serve as pointers to data blocks. These file pointers are called inodes, and each volume has a finite capacity for the number of inodes, which is called the volume file capacity. When a volume runs low on or exhausts its available files (inodes), you can't write additional data to that volume.

The number of file system objects—files, directories, Snapshot copies—a volume can contain is determined by how many inodes it has. The number of inodes in a volume increases commensurately with the volume's storage capacity (and the number of volume constituents for FlexGroup volumes). By default, FlexVol volumes (or FlexGroup constituents) with a storage capacity of 648 GiB or more all have the same number of inodes: 21,251,126. If you create a volume larger than 648 GiB and you want it to have more than 21,251,126 inodes, you must increase the maximum number of inodes (files) manually. For more information about viewing the maximum number of files for a volume, see [Viewing a volume's file capacity](#).

The default number of inodes on a volume is 1 inode for every 32 KiB of volume storage capacity, up to a volume size of 648 GiB. For a 1 GiB volume:

$$\text{Volume_size_in_bytes} \times (1 \text{ file} \div \text{inode_size_in_bytes}) = \text{maximum_number_of_files}$$
$$1,073,741,824 \text{ bytes} \times (1 \text{ file} \div 32,768 \text{ bytes}) = 32,768 \text{ files}$$

You can increase the maximum number of inodes that a volume can contain, up to a maximum of 1 inode for every 4 KiB of storage capacity. For a 1 GiB volume, this increases the maximum number of inodes or files from 32,768 to 262,144:

$$1,073,741,824 \text{ bytes} \times (1 \text{ file} \div 4096 \text{ bytes}) = 262,144 \text{ files}$$

An FSx for ONTAP volume can have a maximum of 2 billion inodes.

For information about changing the maximum number of files that a volume can store, see [Increasing the maximum number of files on a volume](#).

Updating a volume's storage capacity

You can manage volume storage capacity by manually increasing or decreasing volume size using the Amazon Web Services Management Console, Amazon CLI and API, and the ONTAP CLI. You can also enable volume autosizing so that the volume size automatically grows or shrinks when it reaches certain used storage capacity thresholds. You use the ONTAP CLI to manage volume autosizing.

To change a volume's storage capacity (console)

- You can increase or decrease a volume's storage capacity using the Amazon FSx console, Amazon CLI, and API. For more information, see [Updating a volume](#).

You can also use the ONTAP CLI to modify a volume's storage capacity using the [volume modify](#) command.

To modify a volume's size (ONTAP CLI)

- To access the NetApp ONTAP CLI, establish an SSH session on the management port of the Amazon FSx for NetApp ONTAP file system by running the following command. Replace *management_endpoint_ip* with the IP address of the file system's management port.

```
[~]$ ssh fsxadmin@management_endpoint_ip
```

For more information, see [Managing file systems with the ONTAP CLI](#).

- Use the **volume modify** ONTAP CLI command to modify a volume's storage capacity. Run the following command, using your data in place of the following values:
 - Replace *svm_name* with the name of the storage virtual machine (SVM) that the volume is created on.
 - Replace *vol_name* with name of the volume that you want to re-size.
 - Replace *vol_size* with the new size of the volume in the format *integer*[KB|MB|GB|TB|PB]; for example, 100GB to increase the volume size to 100 gigabytes.

```
::> volume modify -server svm_name -volume vol_name -size vol_size
```

Enabling volume autosizing

Volume autosizing so that the volume will automatically grow to a specified size when it reaches a used space threshold. You can do this for FlexVol volume types (the default volume type for FSx for ONTAP) using the [volume autosize](#) ONTAP CLI command.

To enable volume autosizing (ONTAP CLI)

1. To access the NetApp ONTAP CLI, establish an SSH session on the management port of the Amazon FSx for NetApp ONTAP file system by running the following command. Replace *management_endpoint_ip* with the IP address of the file system's management port.

```
[~]$ ssh fsxadmin@management_endpoint_ip
```

For more information, see [Managing file systems with the ONTAP CLI](#).

2. Use the `volume autosize` command as shown, replacing the following values:
 - Replace *svm_name* with the name of the SVM that the volume is created on.
 - Replace *vol_name* with name of the volume that you want to resize.
 - Replace *grow_threshold* with a used space percentage value (such as 90) at which the volume will automatically increase in size (up to the *max_size* value).
 - Replace *max_size* with the maximum size that the volume can grow to. Use the format *integer*[KB|MB|GB|TB|PB]; for example, 300TB. The maximum size is 300 TB. The default is 120% of the volume size.
 - Replace *min_size* with the minimum size that the volume will shrink to. Use the same format as for *max_size*.
 - Replace *shrink_threshold* with the used space percentage at which the volume will automatically shrink in size.

```
::> volume autosize -vserver svm_name -volume vol_name -mode grow_shrink -  
grow-threshold-percent grow_threshold -maximum-size max_size -shrink-threshold-  
percent shrink_threshold -minimum-size min_size
```



Monitoring volume storage capacity

You can view a volume's available storage and its storage distribution in Amazon Web Services Management Console, Amazon CLI, and the NetApp ONTAP CLI.

To monitor a volume's storage capacity (console)

The **Available storage** graph displays the amount of free storage capacity on a volume over time. The **Storage distribution** graph shows how a volume's storage capacity is currently distributed over 4 categories:

- User data
 - Snapshot data
 - Available volume capacity
 - Other data
1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
 2. Choose **Volumes** in the left navigation column, then choose the ONTAP volume that you want to view storage capacity information for. The volume detail page appears.
 3. In the second panel, choose the **Monitoring** tab. The **Available storage** and **Storage distribution** graphs display, along with several other graphs.



To monitor a volume's storage capacity (ONTAP CLI)

You can monitor how your volume's storage capacity is being consumed by using the `volume show-space` ONTAP CLI command. For more information, see [volume show-space](#) in the NetApp ONTAP Documentation Center.

1. To access the NetApp ONTAP CLI, establish an SSH session on the management port of the Amazon FSx for NetApp ONTAP file system by running the following command. Replace *management_endpoint_ip* with the IP address of the file system's management port.

```
[~]$ ssh fsxadmin@management_endpoint_ip
```

For more information, see [Managing file systems with the ONTAP CLI](#).

2. View a volume's storage capacity usage by issuing the following command, replacing the following values:
 - Replace *svm_name* with the name of the SVM that the volume is created on.
 - Replace *vol_name* with name of the volume for which you are setting the data-tiering policy.

```
::> volume show-space -vserver svm_name -volume vol_name
```

If the command is successful, you'll see output similar to the following:

```
Vserver : svm_name
Volume  : vol_name
Feature                               Used           Used%
-----
User Data                             140KB          0%
Filesystem Metadata                   164.4MB        1%
Inodes                                10.28MB        0%
Snapshot Reserve                       563.2MB        5%
Deduplication                          12KB           0%
Snapshot Spill                          9.31GB         85%
Performance Metadata                   668KB          0%
Total Used                             10.03GB        91%
```

Total Physical Used	10.03GB	91%
---------------------	---------	-----

The output of this command shows the amount of physical space that different types of data occupy on this volume. It also shows the percentage of the total volume's capacity that each type of data consumes. In this example, Snapshot `Spill` and Snapshot `Reserve` consume a combined 90 percent of the volume's capacity.

Snapshot `Reserve` shows the amount of disk space reserved for storing Snapshot copies. If the Snapshot copies storage exceeds the reserve space, it spills into the file system and this amount is shown under Snapshot `Spill`.

To increase the amount of available space, you can either [increase the size](#) of the volume, or you can [delete snapshots](#) that you are not using, as shown in the following procedures.

For FlexVol volume types (the default volume type for FSx for ONTAP volumes), you can also enable [volume autosizing](#). When you enable autosizing, the volume size automatically increases when it reaches certain thresholds. You can also disable automatic snapshots. Both of these features are explained in the following sections.

Setting a volume's tiering policy

You can modify a volume's tiering policy using the Amazon Web Services Management Console, Amazon CLI and API, and the ONTAP CLI.

To modify a volume's data tiering policy (console)

Use the following procedure to modify a volume's data-tiering policy using the Amazon Web Services Management Console.

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. Choose **Volumes** in the left navigation pane, then choose the ONTAP volume for which you want to modify the data-tiering policy.
3. Choose **Update volume** from the **Actions** drop down menu. The **Update volume** window appears.
4. For **Capacity pool tiering policy**, choose the new policy for the volume. For more information, see [Volume tiering policies](#).
5. Choose **Update** to apply the new policy to the volume.

To set a volume's tiering policy (CLI)

- Modify a volume's tiering policy using the [update-volume](#) CLI command ([UpdateVolume](#) is the equivalent Amazon FSx API action). The following CLI command example sets a volume's data-tiering policy to SNAPSHOT_ONLY.

```
aws fsx update-volume \  
  --volume-id fsxvol-abcde0123456789f \  
  --ontap-configuration TieringPolicy={Name=SNAPSHOT_ONLY}
```

For a successful request, the system responds with the volume description.

```
{  
  "Volume": {  
    "CreationTime": "2021-10-05T14:27:44.332000-04:00",  
    "FileSystemId": "fs-abcde0123456789f",  
    "Lifecycle": "CREATED",  
    "Name": "vol1",  
    "OntapConfiguration": {  
      "FlexCacheEndpointType": "NONE",  
      "JunctionPath": "/vol1",  
      "SecurityStyle": "UNIX",  
      "SizeInMegabytes": 1048576,  
      "StorageEfficiencyEnabled": true,  
      "StorageVirtualMachineId": "svm-abc0123de456789f",  
      "StorageVirtualMachineRoot": false,  
      "TieringPolicy": {  
        "CoolingPeriod": 2,  
        "Name": "SNAPSHOT_ONLY"  
      },  
      "UUID": "aaaa1111-bb22-cc33-dd44-abcde01234f5",  
      "OntapVolumeType": "RW"  
    },  
    "ResourceARN": "arn:aws:fsx:us-east-2:111122223333:volume/fs-  
abcde0123456789f/fsvol-abc012def3456789a",  
    "VolumeId": "fsvol-abc012def3456789a",  
    "VolumeType": "ONTAP"  
  }  
}
```

To modify a volume's tiering policy (ONTAP CLI)

You use the `volume modify` ONTAP CLI command to set a volume's tiering policy. For more information, see [volume modify](#) in the NetApp ONTAP Documentation Center.

1. To access the NetApp ONTAP CLI, establish an SSH session on the management port of the Amazon FSx for NetApp ONTAP file system by running the following command. Replace *management_endpoint_ip* with the IP address of the file system's management port.

```
[~]$ ssh fsxadmin@management_endpoint_ip
```

For more information, see [Managing file systems with the ONTAP CLI](#).

2. Enter the ONTAP CLI advanced mode using the following command.

```
FSx::> set adv
```

```
Warning: These advanced commands are potentially dangerous; use them only when  
directed to do so by NetApp personnel.
```

```
Do you want to continue? {y|n}: y
```

3. Use the following command to modify the volume data-tiering policy, replacing the following values:
 - Replace *svm_name* with the name of the SVM that the volume is created on.
 - Replace *vol_name* with name of the volume for which you are setting the data-tiering policy.
 - Replace *tiering_policy* with the desired policy. Valid values are snapshot-only, auto, all, or none. For more information, see [Volume tiering policies](#).

```
FSx::> volume modify -vserver svm_name -volume vol_name -tiering-  
policy tiering_policy
```

Setting the minimum cooling days

Minimum cooling days for a volume set the threshold that's used to determine which data is warm and which data is cold. You can set a volume's minimum cooling days using Amazon CLI and API, and the ONTAP CLI.

To set a volume's minimum cooling days (CLI)

- Modify a volume configuration by using the [update-volume](#) CLI command ([UpdateVolume](#) is the equivalent Amazon FSx API action). The following CLI command example sets a volume's `CoolingPeriod` to 104 days.

```
aws fsx update-volume \  
  --volume-id fsxvol-abcde0123456789f \  
  --ontap-configuration TieringPolicy={Name=SNAPSHOT_ONLY} \  
aws fsx update-volume --volume-id fsvol-006530558c14224ac --ontap-configuration \  
  TieringPolicy={CoolingPeriod=104}
```

The system responds with the volume description for a successful request.

```
{  
  "Volume": {  
    "CreationTime": "2021-10-05T14:27:44.332000-04:00",  
    "FileSystemId": "fs-abcde0123456789f",  
    "Lifecycle": "CREATED",  
    "Name": "vol1",  
    "OntapConfiguration": {  
      "FlexCacheEndpointType": "NONE",  
      "JunctionPath": "/vol1",  
      "SecurityStyle": "UNIX",  
      "SizeInMegabytes": 1048576,  
      "StorageEfficiencyEnabled": true,  
      "StorageVirtualMachineId": "svm-abc0123de456789f",  
      "StorageVirtualMachineRoot": false,  
      "TieringPolicy": {  
        "CoolingPeriod": 104,  
        "Name": "SNAPSHOT_ONLY"  
      },  
      "UUID": "aaaa1111-bb22-cc33-dd44-abcde01234f5",  
      "OntapVolumeType": "RW"  
    },  
    "ResourceARN": "arn:aws:fsx:us-east-2:111122223333:volume/fs-  
abcde0123456789f/fsvol-abc012def3456789a",  
    "VolumeId": "fsvol-abc012def3456789a",  
    "VolumeType": "ONTAP"  
  }  
}
```

To set a volume's minimum cooling days (ONTAP CLI)

Use the `volume modify` ONTAP CLI command to set the minimum number of cooling days for an existing volume. For more information, see [volume modify](#) in the NetApp ONTAP Documentation Center.

1. To access the NetApp ONTAP CLI, establish an SSH session on the management port of the Amazon FSx for NetApp ONTAP file system by running the following command. Replace *management_endpoint_ip* with the IP address of the file system's management port.

```
[~]$ ssh fsxadmin@management_endpoint_ip
```

For more information, see [Managing file systems with the ONTAP CLI](#).

2. Enter the ONTAP CLI advanced mode using the following command.

```
FSx::> set adv
```

```
Warning: These advanced commands are potentially dangerous; use them only when  
directed to do so by NetApp personnel.
```

```
Do you want to continue? {y|n}: y
```

3. Use the following command to change your volume's tiering minimum cooling days, replacing the following values:
 - Replace *svm_name* with the name of the SVM that the volume is created on.
 - Replace *vol_name* with name of the volume for which you are setting the cooling days.
 - Replace *cooling_days* with the desired, an integer between 2-183.

```
FSx::> volume modify -vserver svm_name -volume vol_name -tiering-minimum-cooling-  
days cooling_days
```

The system responds as follows for a successful request.

```
Volume modify successful on volume vol_name of Vserver svm_name.
```

Setting a volume's cloud retrieval policy

Use the `volume modify` ONTAP CLI command to set the cloud retrieval policy for an existing volume. For more information, see [volume modify](#) in the NetApp ONTAP Documentation Center.

To set a volume's cloud retrieval policy (ONTAP CLI)

1. To access the NetApp ONTAP CLI, establish an SSH session on the management port of the Amazon FSx for NetApp ONTAP file system by running the following command. Replace *management_endpoint_ip* with the IP address of the file system's management port.

```
[~]$ ssh fsxadmin@management_endpoint_ip
```

For more information, see [Managing file systems with the ONTAP CLI](#).

2. Enter the ONTAP CLI advanced mode using the following command.

```
FSx::> set adv
```

```
Warning: These advanced commands are potentially dangerous; use them only when  
directed to do so by NetApp personnel.  
Do you want to continue? {y|n}: y
```

3. Use the following command to set the volume's cloud retrieval policy, replacing the following values:
 - Replace *svm_name* with the name of the SVM that the volume is created on.
 - Replace *vol_name* with name of the volume for which you are setting the cloud retrieval policy.
 - Replace *retrieval_policy* with the desired value, either `default`, `on-read`, `never`, or `promote`.

```
FSx::> volume modify -vserver svm_name -volume vol_name -cloud-retrieval-  
policy retrieval_policy
```


The system responds as follows for a successful request.

```
Volume modify successful on volume vol_name of Vserver svm_name.
```


Viewing a volume's file capacity

You can use either of the following methods to view the maximum number of files allowed and the number of files already used on a volume.

- The CloudWatch volume metrics `FilesCapacity` and `FilesUsed`.
- In the Amazon FSx console, navigate to the **Available files (inodes)** chart in your volume's **Monitoring** tab. The following image shows the **Available files (inodes)** on a volume decreasing over time.



Increasing the maximum number of files on a volume

FSx for ONTAP volumes can run out of file capacity when the number of available inodes, or file pointers, is exhausted.

To increase the maximum number of files on a volume (ONTAP CLI)

You use the `volume modify` ONTAP CLI command to increase the maximum number of files on a volume. For more information, see [volume modify](#) in the NetApp ONTAP Documentation Center.

1. To access the NetApp ONTAP CLI, establish an SSH session on the management port of the Amazon FSx for NetApp ONTAP file system by running the following command. Replace *management_endpoint_ip* with the IP address of the file system's management port.

```
[~]$ ssh fsxadmin@management_endpoint_ip
```

For more information, see [Managing file systems with the ONTAP CLI](#).

2. Do one of the following, depending on your use case. Replace *svm_name* and *vol_name* with your values.
 - To configure a volume to always have the maximum number of files (inodes) available, perform the following:
 1. Enter advanced mode in the ONTAP CLI by using the following command.

```
::> set adv
```

2. After running this command, you'll see this output. Enter y to continue.

```
Warning: These advanced commands are potentially dangerous; use them only
when
directed to do so by NetApp personnel.
Do you want to continue? {y|n}: y
```

3. Enter the following command to always use the maximum number of files on the volume:

```
::> volume modify -vserver svm_name -volume vol_name -files-set-maximum true
```

- To manually specify the total number of files permitted on the volume, with *max_number_files* = (current_size_of_volume) × (1 file ÷ 4 KiB), up to a maximum possible value of 2 billion, use the following command:

```
::> volume modify -vserver svm_name -volume vol_name -files max_number_files
```

Enabling a volume's cloud write mode

Use the `volume modify` ONTAP CLI command to enable or disable cloud write mode for an existing volume. For more information, see [volume modify](#) in the NetApp ONTAP Documentation Center.

Prerequisites for setting cloud write mode are:

- The volume must be an existing volume. You can only enable the feature on an existing volume.
- The volume must be a read-write (RW) volume.
- The volume must have the **All** tiering policy. For more information about modifying a volume's tiering policy, see [Setting a volume's tiering policy](#).

Cloud write mode is helpful for cases like migrations, for example, where large amounts of data are transferred to a file system using the NFS protocol.

To set a volume's cloud write mode (ONTAP CLI)

1. To access the NetApp ONTAP CLI, establish an SSH session on the management port of the Amazon FSx for NetApp ONTAP file system by running the following command. Replace *management_endpoint_ip* with the IP address of the file system's management port.

```
[~]$ ssh fsxadmin@management_endpoint_ip
```

For more information, see [Managing file systems with the ONTAP CLI](#).

2. Enter the ONTAP CLI advanced mode using the following command.

```
FSx::> set -privilege advanced
Warning: These advanced commands are potentially dangerous; use them only when
        directed to do so by NetApp personnel.
Do you want to continue? {y|n}: y
```

3. Use the following command to set the volume's cloud write mode, replacing the following values:
 - Replace *svm_name* with the name of the SVM that the volume is created on.
 - Replace *vol_name* with name of the volume for which you are setting cloud write mode.

- Replace *vol_cw_mode* with either `true` to enable cloud write mode on the volume or `false` to disable it.

```
FSx::> volume modify -vserver svm_name -volume vol_name -is-cloud-write-  
enabled vol_cw_mode
```

The system responds as follows for a successful request.

```
Volume modify successful on volume vol_name of Vserver svm_name.
```

Protecting your data

Beyond automatically replicating your file system's data to ensure high durability, Amazon FSx provides you with the following options to further protect the data stored on your file systems:

- Native Amazon FSx backups support your backup retention and compliance needs within Amazon FSx. You can also use Amazon Backup to centrally manage, automate, and protect your backups across Amazon Web Services in the cloud.
- Snapshots enable your users to easily undo file changes and compare file versions by restoring files to previous versions.
- Replication of your Amazon FSx file system to a second file system to provide data protection and recovery. Replication, when enabled, occurs on an automatic, scheduled basis.
- SnapLock can protect your files by transitioning them to a write once, read many (WORM) state, which prevents modification or deletion for a specified retention period.

Topics

- [Working with backups](#)
- [Working with snapshots](#)
- [Scheduled replication using NetApp SnapMirror](#)
- [Protecting your data with SnapLock](#)

Working with backups

With FSx for ONTAP, you can take automatic daily backups and user-initiated backups of the volumes on your file system. FSx for ONTAP backups are per volume, so each backup contains only the data in a particular volume. Amazon FSx backups are highly durable and incremental.

All Amazon FSx backups (automatic daily backups and user-initiated backups) are incremental. This 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. When you delete a backup, only the data unique to that backup is removed. Each Amazon FSx backup contains all of the information that is needed to create a new volume from the backup, effectively restoring a point-in-time snapshot of the file system volume.

Creating regular backups for your volumes is a best practice that helps support your data retention and compliance needs. Working with Amazon FSx backups is easy, whether it's creating backups, restoring from a backup, or deleting a backup.

Amazon FSx supports backing up ONTAP FlexVol volumes (on all file systems) and FlexGroup volumes with an `OntapVolumeType` of RW (read-write).

Note

Amazon FSx does not support backing up data protection (DP) volumes, load-sharing (LS) volumes, or destination FlexCache volumes.

There are limits to the number of backups that you can store per file system and per volume. For more information, see [Quotas that you can increase](#) and [Resource quotas for each file system](#).

Topics

- [How backups work](#)
- [Storage requirements](#)
- [Working with automatic daily backups](#)
- [Working with user-initiated backups](#)
- [Copying tags to backups](#)
- [Backup and restore performance](#)
- [Using Amazon Backup with Amazon FSx](#)
- [Restoring backups to a new volume](#)
- [Deleting backups](#)
- [Backups and offline volumes](#)
- [Creating a user-initiated backup](#)
- [Restoring a backup to a new volume](#)
- [Deleting a backup](#)

How backups work

Amazon FSx backups use snapshots – point-in-time, read-only images of your volumes – to maintain incrementality between backups. Each time a backup is taken, Amazon FSx first takes a

snapshot of your volume. The backup snapshot is stored in your volume, and consumes space on your SSD storage tier. Amazon FSx then compares this snapshot to the previous backup snapshot (if one exists) and copies only the changed data into your backup.

If no prior backup snapshot exists, then the entire contents of the most recent backup snapshot is copied into your backup. After the latest backup snapshot is successfully taken, Amazon FSx deletes the previous backup snapshot. The snapshot used for the latest backup remains in your volume until the next backup is taken, when the process repeats. To optimize backup storage costs, ONTAP preserves a volume's storage efficiency savings in its backups.

Amazon FSx can't back up volumes that are offline.

Storage requirements

In order to take backups of your volumes, both your volume and your file system must have enough available SSD storage capacity to store a 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. You can [increase a volume's](#) or [file system's](#) SSD storage at anytime to ensure that your backups won't be interrupted.

Working with automatic daily backups

Automatic daily backups of your file system's volumes are enabled by default when you create a file system. You can enable or disable automatic daily backups for a file system at any time. Automatic daily backups occur during the daily backup window, which is automatically set when you create a file system. You can modify the daily backup window at any time. We recommend that you choose a time of the day for your daily backup that is outside of the normal operating hours for the applications that use your volumes for better backup performance. For more information, see [Backup and restore performance](#).

You can set the retention period for automatic daily backups to between 1 and 90 days in the console when creating a file system or at any time. The default automatic daily backup retention period is 30 days. The service deletes an automatic daily backup once its retention period expires. Using the CLI or API you can set the retention period to between 0 and 90 days; setting it to 0 turns off automatic daily backups.

The daily backup window and the backup retention period are file system-level settings that apply to all of the volumes on your file system. You can use the Amazon FSx console, the Amazon CLI, or the API to change the backup window and backup retention period for your file systems, and to turn automatic daily backups on or off. For more information, see [Updating a file system](#).

You can't create a volume backup if the volume is offline. For more information, see [Backups and offline volumes](#).

Note

Automatic daily backups have a maximum retention period of 90 days, but [user-initiated backups](#) that you create, which include backups created using Amazon Backup, are retained forever unless you or the Amazon Backup service deletes them.

You can manually delete an automatic daily backup using the console, CLI, and API. When you delete a volume, you also delete the automatic daily backups for that volume. Amazon FSx provides the option to create a final backup of a volume before you delete it. The final backup is kept forever, unless you delete it. For more information, see [Deleting backups](#).

Working with user-initiated backups

With Amazon FSx, you can manually take backups of your file system's volumes at any time using the Amazon Web Services Management Console, Amazon CLI, and API. Your user-initiated backups are incremental relative to other backups that may have been created for a volume and are retained forever, unless you delete them. User-initiated backups are retained even after you delete the volume or the file system the backups were created on. You can delete user-initiated backups only by using the Amazon FSx console, API, or CLI. They are never automatically deleted by Amazon FSx. For more information, see [Deleting backups](#).

You cannot create a volume backup if the volume is offline. For more information, see [Backups and offline volumes](#).

Copying tags to backups

When you create or update a volume using the CLI or API, you can enable CopyTagsToBackups to [automatically copy any tags](#) on your volume to its backups. However, if you add any tags while creating a user-initiated backup, including naming a backup when you use the console, the service does *not* copy tags from the volume, even if CopyTagsToBackups is enabled.

Backup and restore performance

A variety of factors can influence the performance of backup and restore operations. Backup and restore operations are background processes, which means they have a lower priority relative to

client IO operations. Client IO operations include NFS, CIFS, and iSCSI data read and writes. All background processes, including backup and restore operations, utilize only the unused portion of your file system's throughput capacity, and can take from a few minutes to a few hours to complete depending on the size of your backup and the amount of unused throughput capacity on your file system.

Other factors that affect backup and restore performance include the storage tier in which your data is stored and the dataset profile. We recommend that you create the first backups of your volumes when most of the data is on SSD storage. Datasets containing mostly small files will typically have lower performance as compared to similarly sized datasets that contain mostly large files. This is because processing large numbers of small files consumes more CPU cycles and network overhead than processing fewer large files.

Generally, you can expect the following backup rates when backing up data stored in the SSD storage tier:

- 750 MBps across several concurrent backups containing mostly large files.
- 100 MBps across several concurrent backups containing mostly small files.

Generally, you can expect the following restore rates:

- 250 MBps across several concurrent restores containing mostly large files.
- 100 MBps across several concurrent restores containing mostly small files.

Using Amazon Backup with Amazon FSx

Amazon Backup is a simple and cost-effective way to protect your data by backing up your Amazon FSx for NetApp ONTAP volumes. Amazon Backup is a unified backup service designed to simplify the creation, restoration, and deletion of backups, while providing improved reporting and auditing. Amazon Backup makes it easier to develop a centralized backup strategy for legal, regulatory, and professional compliance. Amazon Backup also makes protecting your Amazon storage volumes, databases, and file systems simpler by providing a central place where you can do the following:

- Configure and audit the Amazon resources that you want to back up.
- Automate backup scheduling.
- Set retention policies.

- Monitor all recent backup, copy, and restore activity.

Amazon Backup uses the built-in backup functionality of Amazon FSx. Backups created using the Amazon Backup console have the same level of file system consistency and performance, are incremental relative to any other Amazon FSx backups you take of your volume (user-initiated or automatic), and offer the same restore options as backups taken through the Amazon FSx console. If you use Amazon Backup to manage these backups, you gain additional functionality, such as the ability to create scheduled backups as frequently as every hour. You can add an additional layer of defense to protect backups from inadvertent or malicious deletions by storing them in an Amazon Backup Vault.

Backups created by Amazon Backup are considered user-initiated backups, and they count toward the user-initiated backup quota for Amazon FSx. For more information, see [Quotas that you can increase](#). You can view and restore backups created by Amazon Backup in the Amazon FSx console, CLI, and API. However, you can't delete backups created by Amazon Backup in the Amazon FSx console, CLI, or API. For more information, see [Getting started with Amazon Backup](#) in the *Amazon Backup Developer Guide*.

Amazon Backup can't back up volumes that are offline.

Restoring backups to a new volume

You can restore a volume backup to a new volume, effectively restoring a point-in-time snapshot of a volume using the console, CLI, or API.

When restoring a backup, all the data is first written to the SSD storage tier before the service begins tiering data to the capacity pool storage according to the [tiering policy](#) you set for the restored volume. When restoring a backup to a volume with a tiering policy of All, a periodic background process tiers the data to the capacity pool. When restoring a backup to a volume with a tiering policy of Snapshot Only or Auto, data is tiered to the capacity pool if the SSD utilization for the file system is greater than 50%, and the cooling rate is determined by the tiering policy's cooling period.

When you restore a FlexGroup volume backup on a file system that has a different number of high-availability (HA) pairs from the original file system, Amazon FSx might add additional constituent volumes to ensure that the constituents are evenly distributed.

For step-by-step instructions to restore a backup to a new volume, see [Restoring a backup to a new volume](#).

Note

A restored volume always has the same volume style as the original volume. You can't change the volume style when restoring.

Deleting backups

You can delete automatic daily backups and user-initiated backups of your volumes. Deleting a backup is a permanent, unrecoverable action. Any data in a deleted backup is also deleted. Do not delete a backup unless you're sure you won't need that backup again in the future. For instructions describing how to delete backups, see [Deleting a backup](#).

You can't delete backups created by Amazon Backup, which have type **Amazon Backup**, in the Amazon FSx console, CLI, or API. For information on deleting backups created by Amazon Backup, see [Deleting backups](#) in the Amazon Backup Developer Guide.

You cannot delete a volume's backup if the volume is offline. For more information, see [Backups and offline volumes](#).

Important

Do not delete the common snapshot on the volume because it is used to maintain incrementality between your backups. Deleting the common snapshot on the volume will cause the next backup to be of the entire volume rather than just an incremental backup.

Backups and offline volumes

You can't create or delete volume backups if that volume is offline. Use the [volume show](#) ONTAP CLI command to determine a volume's current state and status.

To bring an offline volume back online, use the [volume online](#) ONTAP CLI command as in the following example:

```
::> volume online -volume volume_name -server svm_name
```

```
Volume 'vs1:vol1' is now online.
```

Creating a user-initiated backup

The following procedure describes how use the Amazon FSx console to create a user-initiated backup of a volume.

You cannot create a volume backup if the volume is offline. For more information, see [Backups and offline volumes](#).

To create a user-initiated backup of a volume (console)

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. Navigate to **File systems** and choose the ONTAP file system that you want to back up a volume for.
3. Choose the **Volumes** tab.
4. Choose the volume you want to back up.
5. From **Actions**, choose **Create backup**.
6. In the **Create backup** dialog box that opens, provide a name for your backup. Backup names can be a maximum of 256 Unicode characters, including letters, white space, numbers, and the special characters . + - = _ : /
7. Choose **Create backup**.

You have now created a backup of one of your file system's volumes. You can find a table of all your backups in the Amazon FSx console by choosing **Backups** in the left side navigation. You can search for the name you gave your backup, and the table filters to only show matching results.

When you create a user-initiated backup as this procedure described, it has the type `USER_INITIATED`, and it has the `CREATING` status until it is fully available.

Restoring a backup to a new volume

The following procedures describe how to restore an FSx for ONTAP backup to a new volume using the Amazon Web Services Management Console and Amazon CLI.

To restore a volume backup to a new volume (Console)

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. In the navigation pane, choose **Backups**, and then choose the FSx for ONTAP volume backup that you want to restore.

3. In the upper right **Actions** menu, choose **Restore backup**. The **Create volume from backup page** appears.
4. Choose the FSx for ONTAP **File system** and **Storage virtual machine** that you want to restore the backup to from the dropdown menus.
5. Under **Volume details**, there are several selections. First, enter the **Volume name**. You can use a up to 203 alphanumeric or underscore (_) characters.
6. For **Volume size**, enter any whole number in the range of 20–314572800 to specify the size in mebibytes (MiB).
7. For **Volume type**, choose **Read-Write (RW)** to create a volume that is readable and writable or **Data Protection (DP)** to create a volume that is read-only and can be used as the destination of a NetApp SnapMirror or SnapVault relationship. For more information, see [Volume types](#).
8. For **Junction path**, enter a location within the file system to mount the volume. The name must have a leading forward slash, for example /vol3.
9. For **Storage efficiency**, choose **Enabled** to enable the ONTAP storage-efficiency features (deduplication, compression, and compaction). For more information, see [FSx for ONTAP storage efficiency](#).
10. For **Volume security style**, choose either **Unix (Linux)**, **NTFS**, or **Mixed**. A volume's security style determines whether preference is given to NTFS or UNIX ACLs for multi-protocol access. The MIXED mode is not required for multi-protocol access and is only recommended for advanced users.
11. For **Snapshot policy**, choose a snapshot policy for the volume. For more information about snapshot policies, see [Snapshot policies](#).

If you choose **Custom policy**, you must specify the policy's name in the **custom-policy** field. The custom policy must already exist on the SVM or in the file system. You can create a custom snapshot policy with the ONTAP CLI or REST API. For more information, see [Create a Snapshot Policy](#) in the NetApp ONTAP Product Documentation.

12. For **Tiering policy cooling period**, valid values are 2-183 days. A volume's tiering policy cooling period defines the number of days before data that has not been accessed is marked cold and moved to capacity pool storage. This setting only affects the Auto and Snapshot-only policies.
13. In the **Advanced** section, for **SnapLock Configuration**, you can leave the default **Disabled** setting or choose **Enabled** to configure a SnapLock volume. For more information about configuring a SnapLock Compliance volume or a SnapLock Enterprise volume, see [Creating](#)

[a SnapLock Compliance volume](#) and [Creating a SnapLock Enterprise volume](#). For more information about SnapLock, see [Protecting your data with SnapLock](#).

14. Choose **Confirm** to create the volume.

To restore a volume backup to a new volume (CLI)

Use the [create-volume-from-backup](#) CLI command, or the equivalent [CreateVolumeFromBackup](#) API command to restore a volume backup to a new volume.

```
$ aws fsx create-volume-from-backup --backup-id backup-08e6fc1133fff3532 \
  --name demo --ontap-configuration JunctionPath=/demo, SizeInMegabytes=100000, \
  StorageVirtualMachineId=svm-0f04a9c7c27e1908b, TieringPolicy={Name=ALL}
```

The system response for a successful request:

```
{
  "Volume": {
    "CreationTime": 1692721488.428,
    "FileSystemId": "fs-07ab735385276ed60",
    "Lifecycle": "CREATING",
    "Name": "demo",
    "OntapConfiguration": {
      "FlexCacheEndpointType": "NONE",
      "JunctionPath": "/demo",
      "SizeInMegabytes": 100000,
      "StorageEfficiencyEnabled": true,
      "StorageVirtualMachineId": "svm-0f04a9c7c27e1908b",
      "StorageVirtualMachineRoot": false,
      "TieringPolicy": {
        "Name": "ALL"
      },
      "OntapVolumeType": "DP",
      "SnapshotPolicy": "default",
      "CopyTagsToBackups": false,
    },
    "ResourceARN": "arn:aws:fsx:us-east-1:752825163408:volume/fs-07ab735385276ed60/fsvol-0b6ec764c9c5f654a",
    "VolumeId": "fsvol-0b6ec764c9c5f654a",
    "VolumeType": "ONTAP",
  }
}
```

```
}  
}
```

Deleting a backup

You can delete automatic daily backups and user-initiated backups using the Amazon FSx console, CLI, and API, as described in the following procedures.

To delete backups created using Amazon Backup, see [Deleting backups](#) in the Amazon Backup Developer Guide.

To delete a backup (console)

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. From the console dashboard, choose **Backups** from the left side navigation.
3. Choose the backup that you want to delete from the **Backups** table, and then choose **Delete backup**.
4. In the **Delete backups** dialog box that opens, confirm that the ID of the backup shown is the backup that you want to delete.
5. Confirm that the check box is checked for the backup that you want to delete.
6. Choose **Delete backups**.

Your backup and all included data are now permanently and irrecoverably deleted.

To delete a backup (CLI)

- Use the delete-backup CLI command or the equivalent DeleteBackup API action to delete an FSx for ONTAP volume backup, as shown in the following example.

```
$ aws fsx delete-backup --backup-id backup-a0123456789abcdef
```

The system response includes the ID of the backup being deleted, and its lifecycle status, with DELETED indicating that the request was successful.

```
{  
  "BackupId": "backup-a0123456789abcdef",  
  "Lifecycle": "DELETED"
```

}

Working with snapshots

A *snapshot* is a read-only image of an Amazon FSx for NetApp ONTAP volume at a point in time. Snapshots offer protection against accidental deletion or modification of files in your volumes. With snapshots, your users can easily view and restore individual files or folders from an earlier snapshot. Doing this enables users to easily undo changes and compare file versions.

Because snapshots are stored alongside your file system's data, they consume the file system's storage capacity. However, snapshots consume storage capacity only for the changed portions of files since the last snapshot. Note that snapshots consume no capacity at the time of creation. Snapshots stored in your file system are not included in backups of your file system volumes.

Snapshots are enabled by default on your volumes using the default snapshot policy. Your FSx for ONTAP file system has three built-in snapshot policies to choose from when creating or updating a volume in the Amazon FSx console, the Amazon CLI, or the Amazon FSx API. You can also create a custom snapshot policy or snapshots on-demand using the ONTAP CLI or REST API. Snapshots are stored in the `.snapshot` directory at the root of a volume. You can store up to 1,023 snapshots per volume at any point in time. Once you reach this limit, you must delete an existing snapshot before a new snapshot of your volume can be created.

Topics

- [Snapshot policies](#)
- [Restoring individual files and folders](#)
- [Restore files from Snapshots](#)
- [Deleting snapshots](#)
- [Create a Snapshot autodelete policy](#)
- [Delete snapshots](#)
- [Disabling automatic snapshots](#)
- [Snapshot reserve](#)
- [Updating volume's Snapshot reserve](#)

Snapshot policies

The snapshot policy defines how the system creates snapshots for a volume. The policy specifies when to create snapshots, how many copies to retain, and how to name them. There are 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.

The `default` 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.

Note

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, see [Displaying and setting the system time zone](#) in the NetApp Support documentation.

The `default-1weekly` policy works in the same way as the `default` policy, except that it only retains one snapshot from the weekly schedule.

The `none` policy doesn't take any snapshots. You can assign this policy to volumes to prevent automatic snapshots from being taken.

You can also create a custom snapshot policy using the ONTAP CLI or REST API. For more information, see [Create a Snapshot Policy](#) in the *NetApp ONTAP Product Documentation*. You can choose a snapshot policy while creating or updating a volume in the Amazon FSx console, the

Amazon CLI, or the Amazon FSx API. For more information, see [Creating volumes](#) and [Updating a volume](#).

Restoring individual files and folders

Using the snapshots on your Amazon FSx file system, your users can quickly restore previous versions of individual files or folders. Doing this enables them to recover deleted or changed files stored on the shared file system. They do this in a self-service manner directly on their desktop without administrator assistance. This self-service approach increases productivity and reduces administrative workload.

Linux and macOS clients can view snapshots in the `.snapshot` directory at the root of a volume. Windows clients can view snapshots in the `Previous Versions` tab of Windows Explorer (when right-clicking on a file or folder).

Restore files from Snapshots

To restore a file from a snapshot (Linux and macOS clients)

1. If the original file still exists and you do not want it overwritten by the file in a snapshot, then use your Linux or macOS client to rename the original file or move it to a different directory.
2. In the `.snapshot` directory, locate the snapshot that contains the version of the file that you want to restore.
3. Copy the file from the `.snapshot` directory to the directory in which the file originally existed.

To restore a file from a snapshot (Windows clients)

Users on Windows clients can restore files to previous versions using the familiar Windows File Explorer interface.

1. To restore a file, users choose the file to restore, then choose **Restore previous versions** from the context (right-click) menu.
2. Users can then view and restore a previous version from the **Previous Versions** list.

Data in snapshots is read-only. If you want to make modifications to files and folders listed in the **Previous Versions** tab, you must save a copy of the files and folders that you want to modify to a writable location and make modifications to the copies.

Deleting snapshots

[Snapshots](#) are point-in-time read-only images of a prior state of your volume, and are enabled by default on all FSx for ONTAP volumes to protect your data. Snapshots consume storage capacity only for the portions of files that have changed since the last snapshot. For this reason, if your workload changes data rapidly, snapshots from old data can take up a significant portion of the capacity of your volume.

For example, the `volume show-space` command output provided earlier shows 140 KB of User Data. However, the volume had 9.8 GB of User Data before the user data was deleted. Even if you've deleted the files from your volume, a snapshot might still reference old user data. Because of this, `Snapshot Reserve` and `Snapshot Spill` in the prior example take up a total of 9.8 GB of space, even though there is virtually no user data on the volume.

To free up space on volumes, you can delete older snapshots that you no longer need. You can do this by creating a snapshot auto-delete policy or by manually deleting snapshots. Deleting a snapshot deletes the changed data stored on the snapshot.

Create a Snapshot autodelete policy

You can create a policy to automatically delete snapshots when the amount of available space in your volume is running low. To establish an autodelete policy for a volume, use the following command.

Before running this command, replace the following values:

- Replace *svm_name* with the name of the SVM that the volume is created on.
- Replace *vol_name* with name of the volume.

For `-trigger`, assign one of the following values:

- `volume` – Use `volume` if you want the threshold at which snapshots are deleted to correspond to a total used-volume capacity threshold. The used-volume capacity thresholds that trigger snapshot deletion are determined by the size of your volume, with the threshold scaling from 85–98 percent used capacity. Smaller volumes have a smaller threshold, and larger volumes have a larger one.
- `snap_reserve` – Use `snap_reserve` if you want snapshots to be deleted based on what can be held in your snapshot reserve.

```
::> volume snapshot autodelete modify -vserver svm_name -volume vol_name -enabled true  
-trigger [volume|snap_reserve]
```

For more information, see the [volume snapshot autodelete modify](#) command in the *NetApp ONTAP Documentation Center*.

Delete snapshots

Use the following command to manually delete snapshots. Before running this command, replace the following values:

- Replace *svm_name* with the name of the SVM that the volume is created on.
- Replace *vol_name* with name of the volume.
- Replace *snapshot_name* with the name of the snapshot. This command supports wildcard characters (*) for *snapshot_name*. Therefore, you can delete all hourly snapshots, for example, by using `hourly*`.

Important

If you have Amazon FSx backups enabled, Amazon FSx retains a snapshot for the most recent Amazon FSx backup of each volume. Those snapshots are used to maintain incrementality between backups, and must not be deleted by using this method.

```
FsxIdabcdef01234567892::> volume snapshot delete -vserver svm_name -volume vol_name -  
snapshot snapshot_name
```

For more information about manually deleting snapshots, see the [volume snapshot delete](#) command in the *NetApp ONTAP Documentation Center*.

Disabling automatic snapshots

Automatic snapshots are enabled by the default snapshot policy for volumes in your FSx for ONTAP file system. If you don't need snapshots of your data (for example, if you're using test data), you can disable snapshots by setting the volume's [snapshot policy](#) to none using the Amazon Web Services Management Console, Amazon CLI and API, and the ONTAP CLI, as described in the following procedures.

To disable automatic snapshots (Amazon Web Services Management Console)

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. Navigate to **File systems** and choose the ONTAP file system that you want to update a volume for.
3. Choose the **Volumes** tab.
4. Choose the volume that you want to update.
5. For **Actions**, choose **Update volume**.

The **Update volume** dialog box displays with the volume's current settings.

6. For **Snapshot policy**, choose **None**.
7. Choose **Update** to update the volume.

To disable automatic snapshots (Amazon CLI)

- Use the [update-volume](#) CLI command (or the equivalent [UpdateVolume](#) API operation), to set the SnapshotPolicy to none, as shown in the following example.

```
aws fsx update-volume \
  --volume-id fsvol-1234567890abcdefa \
  --name new_vol \
  --ontap-configuration CopyTagsToBackups=true,JunctionPath=/new_vol, \
    SizeInMegabytes=2048,SnapshotPolicy=none, \
    StorageEfficiencyEnabled=true, \
    TieringPolicy=all
```

To disable automatic snapshots (ONTAP CLI)

1. To show the none policy, use the following command.

```
::> snapshot policy show -policy none

Vserver: FsxIdabcdef01234567892
Policy Name          Number of Is Schedules Enabled Comment
-----
none                 0 false      Policy for no automatic snapshots.
Schedule            Count      Prefix      SnapMirror Label
```

```
-----  
- - - - -
```

2. To disable automatic snapshots, add the none policy to your volume by using the following command.
 - Replace *svm_name* with your SVM's name.
 - Replace *vol_name* with your volume's name.

When prompted to continue, enter **y**.

```
::> volume modify -vserver svm_name -volume vol_name -snapshot-policy none
```

```
Warning: You are changing the Snapshot policy on volume "vol_name" to "none".  
Snapshot copies on this volume  
    that do not match any of the prefixes of the new Snapshot policy will not  
be deleted. However, when  
    the new Snapshot policy takes effect, depending on the new retention  
count, any existing Snapshot copies  
    that continue to use the same prefixes might be deleted. See the 'volume  
modify' man page for more information.  
Do you want to continue? {y|n}: y  
Volume modify successful on volume vol_name of Vserver svm_name.
```

Snapshot reserve

Snapshot copy reserve sets a specific percent of the disk space for storing Snapshot copies. The default Snapshot copy reserve is set to 5 percent of the disk space. If the Snapshot copies exceed the reserve space, they spill into the active file system and this process is called Snapshot spill.

The Snapshot copy reserve must have sufficient space allocated for the Snapshot copies, including [volume backups](#). If the Snapshot copies exceeds the reserve space, you must delete the existing Snapshot copies from the active file system to recover the space, for the use of the file system. You can also modify the percent of disk space that is allotted to Snapshot copies.

Whenever Snapshots consume more than 100% of the Snapshot reserve, they begin to occupy primary SSD storage space. This process is called Snapshot spill. When the Snapshots continue to occupy the active file system space, the system is in danger of becoming full. If the system becomes full due to Snapshot spill, you can create files only after you delete enough Snapshots.

When enough disk space is available for Snapshots in the Snapshot reserve, deleting files from the primary SSD tier frees disk space for new files, while the Snapshot copies that reference those files consume only the space in the Snapshot copy reserve.

Because there is no way to prevent Snapshots from consuming disk space greater than the amount reserved for them (the Snapshot reserve), it is important to reserve enough disk space for Snapshots so that the primary SSD tier always has space available to create new files or modify existing ones.

If a Snapshot is created when the disks are full, deleting files from the primary SSD tier does not create any free space because all that data is also referenced by the newly created Snapshot. You must [delete the Snapshot](#) before in order to free up storage in order to create or update any files.

You can modify the amount of Snapshot reserve on a volume using the NetApp ONTAP CLI. For more information, see [Updating volume's Snapshot reserve](#).

Updating volume's Snapshot reserve

You can change the amount of Snapshot reserve on a volume using the NetApp ONTAP CLI or API, described in the following procedure.

1. To access the NetApp ONTAP CLI, establish an SSH session on the management port of the Amazon FSx for NetApp ONTAP file system by running the following command. Replace *management_endpoint_ip* with the IP address of the file system's management port.

```
[~]$ ssh fsxadmin@management_endpoint_ip
```

For more information, see [Managing file systems with the ONTAP CLI](#).

2. Use the `snap reserve` ONTAP CLI command to change the percent of disk space used for the Snapshot copy reserve. Replace *vol_name* with the name of the volume, and *percent* with the percent of disk space you want to reserve for Snapshot copies.

```
::> snap reserve vol_name percent
```

The following example changes the snapshot reserve for `vol1` to 25% of the volumes storage capacity.

```
::> snap reserve vol1 25
```

Scheduled replication using NetApp SnapMirror

You can use NetApp SnapMirror to schedule periodic replication of your FSx for ONTAP file system to or from a second file system. This capability is available for both in-Region and cross-Region deployments.

NetApp SnapMirror replicates data at high speeds, so you get high data availability and fast data replication across ONTAP systems, whether you're replicating between two Amazon FSx file systems in Amazon, or from on-premises to Amazon. Replication can be scheduled as frequently as every 5 minutes, although intervals should be carefully chosen based on RPOs (Recovery Point Objectives), RTOs (Recovery Time Objectives), and performance considerations.

When you replicate data to NetApp storage systems and continually update the secondary data, your data is kept current and remains available whenever you need it. No external replication servers are required. For more information about using NetApp SnapMirror to replicate your data, see [Learn about the Replication service](#) in the *NetApp BlueXP documentation*.

You can create a data protection (DP) destination volume for NetApp SnapMirror using the Amazon FSx console, the Amazon CLI, and the Amazon FSx API, in addition to the NetApp ONTAP CLI and REST API. For information about creating a destination volume using the Amazon FSx console and Amazon CLI, see [Creating volumes](#).

You can use NetApp BlueXP or the NetApp ONTAP CLI to schedule replication for your file system.

Note

There are two types of SnapMirror replication: Volume-level SnapMirror and SVM Disaster Recovery (SVMDR). Only volume-level SnapMirror replication is supported by FSx for ONTAP.

Using NetApp BlueXP to schedule replication

You can use NetApp BlueXP to set up replication with SnapMirror on your FSx for ONTAP file system. For more information, see [Replicating data between systems](#) in the *NetApp BlueXP documentation*.

Using the NetApp ONTAP CLI to schedule replication

You can use the NetApp ONTAP CLI to configure scheduled volume replication. For information, see [Managing SnapMirror volume replication](#) in the *NetApp ONTAP Documentation Center*.

Protecting your data with SnapLock

SnapLock is a feature that allows you to protect your files by transitioning them to a write once, read many (WORM) state, which prevents modification or deletion for a specified retention period. You can use SnapLock to meet regulatory compliance, to protect business-critical data from ransomware attacks, and to provide an additional layer of protection for your data against alteration or deletion.

Amazon FSx for NetApp ONTAP supports the Compliance and Enterprise modes of retention with SnapLock. For more information, see [SnapLock Compliance](#) and [SnapLock Enterprise](#).

You can create SnapLock volumes on FSx for ONTAP file systems created on or after July 13, 2023. Existing file systems will get SnapLock support during an upcoming weekly maintenance window.

Topics

- [How SnapLock works](#)
- [SnapLock Compliance](#)
- [SnapLock Enterprise](#)
- [Working with the retention period in SnapLock](#)
- [Committing files to WORM state](#)
- [Backing up SnapLock volumes](#)
- [Deleting SnapLock volumes](#)

How SnapLock works

SnapLock can help you meet regulatory and governance purposes by preventing your files from being deleted, changed, or renamed. When you create a SnapLock volume, you commit your files to write once, read many (WORM) storage and set retention periods for the data. Your files can be stored in a non-erasable, non-writable state for a designated period, or indefinitely.

Important

You must specify whether a volume will use SnapLock settings at the time of creation. A non-SnapLock volume can't be converted to a SnapLock volume after creation.

Retention modes

SnapLock has two retention modes: Compliance and Enterprise. Amazon FSx for NetApp ONTAP supports both of them. They have different use cases and some of the features differ, but they both protect your data from modification or deletion using the WORM model. The following table explains some of the similarities and differences between these retention modes.

SnapLock feature	SnapLock Compliance	SnapLock Enterprise
Description	Files transitioned to WORM on a Compliance volume can't be deleted until their retention periods expire.	Files transitioned to WORM on an Enterprise volume can be deleted by authorized users before their retention periods expire using privileged delete.
Use cases	<ul style="list-style-type: none"> To address government or industry-specific mandates such as SEC Rule 17a-4(f), FINRA Rule 4511, and CFTC Regulation 1.31. To protect against ransomware attacks. 	<ul style="list-style-type: none"> To advance an organization's data integrity and internal compliance. To test retention settings before using SnapLock Compliance.
Autocommit	Yes	Yes
Event-based retention (EBR)*	Yes	Yes
Legal Hold*	Yes	No
Privileged delete	No	Yes

SnapLock feature	<u>SnapLock Compliance</u>	<u>SnapLock Enterprise</u>
<u>Volume-append mode</u>	Yes	Yes
<u>SnapLock audit log volumes</u>	Yes	Yes

* EBR and Legal Hold operations are supported in the ONTAP CLI and REST API.

SnapLock administrator

You must have SnapLock administrator privileges to perform certain actions on SnapLock volumes. SnapLock administrator privileges are defined in the `vsadmin-snaplock` role in the ONTAP CLI. You must be a cluster administrator to create a storage virtual machine (SVM) administrator account with the SnapLock administrator role.

You can perform the following actions with the `vsadmin-snaplock` role in the ONTAP CLI:

- Manage your own user account, local password, and key information
- Manage volumes, except moving volumes
- Manage quotas, qtrees, snapshot copies, and files
- Perform SnapLock actions, including privileged delete and Legal Hold
- Configure Network File System (NFS) and Server Message Block (SMB) protocols
- Configure Domain Name System (DNS), Lightweight Directory Access Protocol (LDAP), and Network Information Service (NIS) services
- Monitor jobs

The following procedure details how to create a SnapLock administrator in the ONTAP CLI. You must be logged in as a cluster administrator on a secure connection, such as Secure Shell Protocol (SSH) to perform this task.

To create an SVM administrator account with the `vsadmin-snaplock` role in the ONTAP CLI

- Run the following command. Replace *SVM_name* and *SnapLockAdmin* with your own information.

```
cluster1::> security login create -vserver SVM_name -user-or-group-  
name SnapLockAdmin -application ssh -authentication-method password -role vsadmin-  
snaplock
```

SnapLock audit log volumes

A SnapLock audit log volume contains SnapLock audit logs, which contain timestamps of events such as when a SnapLock administrator was created, when privileged delete operations were executed, or when a Legal Hold was placed on files. The SnapLock audit log volume is a non-erasable record of events.

You must create a SnapLock audit log volume in the same SVM as the SnapLock volume for the following actions:

- To turn on or turn off privileged delete on a SnapLock Enterprise volume.
- To apply Legal Hold on a file in a SnapLock Compliance volume.

Warning

- The minimum retention period for a SnapLock audit log volume is six months. Until this retention period expires, the SnapLock audit log volume and the SVM and file system that are associated with it can't be deleted even if the volume was created in SnapLock Enterprise mode.
- If a file is deleted using privileged delete and its retention period is longer than the retention period of the volume, then the audit log volume inherits the file's retention period. For example, if a file that has a retention period of 10 months is deleted using privileged delete and the retention period of the audit log volume is six months, the retention period of the audit log volume is extended to 10 months.

You can have only one active SnapLock audit log volume in an SVM, but it can be shared by multiple SnapLock volumes in the SVM. To mount a SnapLock audit log volume successfully, set the junction path to `/snaplock_audit_log`. No other volumes can use this junction path, including volumes that aren't audit log volumes.

You can find SnapLock audit logs in the `/snaplock_log` directory under the root of the audit log volume. Privileged delete operations are logged in the `privdel_log` subdirectory. Legal Hold begin and end operations are logged in `/snaplock_log/legal_hold_logs/`. All other logs are stored in the `system_log` subdirectory.

You can create a SnapLock audit log volume with the Amazon FSx console, the Amazon CLI, the Amazon FSx API, and the ONTAP CLI and REST API.

Note

A data protection (DP) volume can't be used as a SnapLock audit log volume.

The following procedure explains how to create a SnapLock audit log volume on the Amazon FSx console.

To create a SnapLock audit log volume the Amazon FSx console

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. Follow the procedure for creating a new volume in [Creating volumes](#).
3. In the **Advanced** section, for **SnapLock Configuration**, choose **Enabled**.

Select the check box to acknowledge the warning about enabling SnapLock on the volume.

4. For **Audit log volume**, choose **Enabled**.

Make sure that the **Junction path** is set to `/snaplock_audit_log`.

5. Follow the rest of the procedure for creating a new volume in [Creating volumes](#).
6. Choose **Confirm** to create the volume.

To turn on the SnapLock audit log volume with the Amazon FSx API, use `AuditLogVolume` in the [CreateSnaplockConfiguration](#).

Accessing your data in a SnapLock volume

You can use open file protocols such as NFS and SMB to access your data in a SnapLock volume. There is no performance impact from writing data to a SnapLock volume or from reading data that's protected by WORM.

You can copy files across SnapLock volumes with NFS and SMB, but they won't retain their WORM properties on the destination SnapLock volume. You must recommit the copied files to WORM to prevent them from being modified or deleted. For more information, see [Committing files to WORM state](#).

You can also replicate SnapLock data with SnapMirror, but the source and destination volumes must be SnapLock volumes with the same retention mode (for example, both must be Compliance or Enterprise).

SnapLock Compliance

Amazon FSx for NetApp ONTAP supports SnapLock Compliance volumes.

Using SnapLock Compliance

This section describes use cases and considerations for the Compliance retention mode.

Use cases for SnapLock Compliance

You might choose the Compliance retention mode for the following use cases.

- You can use SnapLock Compliance to address government or industry-specific mandates such as SEC Rule 17a-4(f), FINRA Rule 4511, and CFTC Regulation 1.31. SnapLock Compliance on Amazon FSx for NetApp ONTAP was assessed for these mandates and regulations by Cohasset Associates. For more information, see the [Compliance Assessment Report for Amazon FSx for NetApp ONTAP](#).
- You can use SnapLock Compliance to complement or enhance a comprehensive data protection strategy to combat ransomware attacks.

Considerations for SnapLock Compliance

Here are some important items to consider about the Compliance retention mode.

- After a file is transitioned to the write once, read many (WORM) state on a SnapLock Compliance volume, it can't be deleted before its retention period expires by any user.
- A SnapLock Compliance volume can only be deleted when the retention periods of all WORM files on the volume have expired, and the WORM files have been deleted from the volume.
- You can't rename a SnapLock Compliance volume after creation.

- You can use SnapMirror to replicate WORM files, but the source volume and destination volume must have the same retention mode (for example, both must be Compliance).
- A SnapLock Compliance volume can't be converted to a SnapLock Enterprise volume, and the reverse.

Creating a SnapLock Compliance volume

You can create a SnapLock Compliance volume with the Amazon FSx console, the Amazon CLI, the Amazon FSx API, and the ONTAP CLI and REST API.

To create a SnapLock compliance volume with the Amazon FSx API, use `SnapLockType` in the [CreateSnaplockConfiguration](#).

The following procedure explains how to create a SnapLock Compliance volume on the Amazon FSx console.

To create a SnapLock Compliance volume on the Amazon FSx console

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. Follow the procedure for creating a new volume in [Creating volumes](#).
3. In the **Advanced** section, for **SnapLock Configuration**, choose **Enabled**.

Select the check box to acknowledge the warning about enabling SnapLock on the volume.

4. For **Retention mode**, choose **Compliance**.
5. For **Audit log volume**, choose between **Enabled** and **Disabled**.

If you choose **Enabled**, make sure that the **Junction path** is set to `/snaplock_audit_log`.

For more information, see [SnapLock audit log volumes](#).

6. For **Retention period**, enter values for **Default retention**, **Minimum retention**, and **Maximum retention**. Then choose a corresponding **Unit** for each.

For more information, see [Working with the retention period in SnapLock](#).

7. For **Autocommit**, choose between **Enabled** and **Disabled**.

If you choose **Enabled**, for **Autocommit period**, enter a value and choose a corresponding **Autocommit unit**.

You can specify a value between 5 minutes and 10 years.

For more information, see [Autocommit](#).

8. For **Volume append mode**, choose between **Enabled** and **Disabled**.

For more information, see [Volume-append mode](#).

9. Follow the rest of the procedure for creating a new volume in [Creating volumes](#).
10. Choose **Confirm** to create the volume.

SnapLock Enterprise

Amazon FSx for NetApp ONTAP supports SnapLock Enterprise volumes.

Using SnapLock Enterprise

This section describes use cases and considerations for the Enterprise retention mode.

Use cases for SnapLock Enterprise

You might choose the Enterprise retention mode for the following use cases.

- You can use SnapLock Enterprise to authorize only specific users to delete files.
- You can use SnapLock Enterprise to advance your organization's data integrity and internal compliance.
- You can use SnapLock Enterprise to test retention settings before using SnapLock Compliance.

Considerations for using SnapLock Enterprise

Here are some important items to consider about the Enterprise retention mode.

- You can use SnapMirror to replicate WORM files, but the source volume and destination volume must have the same retention mode (for example, both must be Enterprise).
- A SnapLock volume can't be converted from Enterprise to Compliance, or from Compliance to Enterprise.
- SnapLock Enterprise doesn't support Legal Hold.

Privileged delete

One of the key differences between SnapLock Enterprise and SnapLock Compliance is that 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. The SnapLock administrator is the only user who can delete files from a SnapLock Enterprise volume that has active retention policies placed on it. For more information, see [SnapLock administrator](#).

You can turn on or turn off privileged delete with the Amazon FSx console, the Amazon CLI, the Amazon FSx API, and the ONTAP CLI and REST API. To turn on privileged delete, you must first create a SnapLock audit log volume in the same SVM as the SnapLock volume. For more information, see [SnapLock audit log volumes](#).

To turn on privileged delete with the Amazon FSx API, use `PrivilegedDelete` in the [CreateSnaplockConfiguration](#).

The following procedure explains how to turn on privileged delete on the Amazon FSx console.

To turn on privileged delete on a SnapLock Enterprise volume on the Amazon FSx console

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. Follow the procedure for creating a new volume in [Creating volumes](#).
3. In the **Advanced** section, for **SnapLock Configuration**, choose **Enabled**.

Select the check box to acknowledge the warning about enabling SnapLock on the volume.

4. For **Retention mode**, choose **Enterprise**.
5. For **Privileged Delete**, choose **Enabled**.
6. Follow the rest of the procedure for creating a new volume in [Creating volumes](#).
7. Choose **Confirm** to create the volume.

Note

You can't issue a privileged delete command to delete a write once, read many (WORM) file that has an expired retention period. You can issue a normal delete operation after the retention period expires.

You can opt to turn off privileged delete permanently, but this action is irreversible. If privileged delete is permanently turned off, you don't need to have a SnapLock audit log volume associated with the SnapLock Enterprise volume.

To permanently turn off privileged delete with the Amazon FSx API, use `PrivilegedDelete` in the [CreateSnaplockConfiguration](#).

To permanently turn off privileged delete on a SnapLock Enterprise volume on the Amazon FSx console

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. Follow the procedure for creating a new volume in [Creating volumes](#).
3. In the **Advanced** section, for **SnapLock Configuration**, choose **Enabled**.

Select the check box to acknowledge the warning about enabling SnapLock on the volume.

4. For **Retention mode**, choose **Enterprise**.
5. For **Privileged Delete**, choose **Permanently disabled**.
6. Follow the rest of the procedure for creating a new volume in [Creating volumes](#).
7. Choose **Confirm** to create the volume.

Creating a SnapLock Enterprise volume

You can create a SnapLock Enterprise volume with the Amazon FSx console, the Amazon CLI, the Amazon FSx API, and the ONTAP CLI and REST API.

To create a SnapLock enterprise volume with the Amazon FSx API, use `SnaplockType` in the [CreateSnaplockConfiguration](#).

To create a SnapLock Enterprise volume on the Amazon FSx console

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. Follow the procedure for creating a new volume in [Creating volumes](#).
3. In the **Advanced** section, for **SnapLock Configuration**, choose **Enabled**.

Select the check box to acknowledge the warning about enabling SnapLock on the volume.

4. For **Retention mode**, choose **Enterprise**.

5. For **Audit log volume**, choose between **Enabled** and **Disabled**.

If you choose **Enabled**, make sure that the **Junction path** is set to `/snaplock_audit_log`.

For more information, see [SnapLock audit log volumes](#).

6. For **Retention period**, enter values for **Default retention**, **Minimum retention**, and **Maximum retention**. Then choose a corresponding **Unit** for each.

For more information, see [Working with the retention period in SnapLock](#).

7. For **Autocommit**, choose between **Enabled** and **Disabled**.

If you choose **Enabled**, for **Autocommit period**, enter a value and choose a corresponding **Autocommit unit**.

You can specify a value between 5 minutes and 10 years.

For more information, see [Autocommit](#).

8. For **Privileged Delete**, choose between **Enabled**, **Disabled**, and **Permanently disabled**.

For more information, see [Privileged delete](#).

9. For **Volume append mode**, choose between **Enabled** and **Disabled**.

For more information, see [Volume-append mode](#).

10. Follow the rest of the procedure for creating a new volume in [Creating volumes](#).

11. Choose **Confirm** to create the volume.

Bypassing Enterprise mode

If you are using the Amazon FSx console or Amazon FSx API, you must have the IAM `fsx:BypassSnapLockEnterpriseRetention` permission to delete a SnapLock Enterprise volume that contains WORM files with active retention policies.

For more information, see [Deleting SnapLock volumes](#).

Working with the retention period in SnapLock

When you create a SnapLock volume, you can set a default retention period for the volume, or you can set the retention period for write once, read many (WORM) files explicitly. During the retention period, you can't delete or modify WORM-protected files. The retention period is used to calculate

the retention time. For example, if you transition a file to WORM on July 14, 2023 at midnight and set the retention period to five years, then the retention time would be until July 14, 2028 at midnight.

For more information about WORM, see [Committing files to WORM state](#).

Retention period policies

The retention period is determined by values that you assign to the following parameters:

- Default retention – The default retention period that's assigned to a WORM file if you don't provide an explicit retention period for it.
- Minimum retention – The shortest retention period that can be assigned to a WORM file.
- Maximum retention – The longest retention period that can be assigned to a WORM file.

Note

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.

You can specify the retention period using several different units of time. The following table lists the specific ranges that are supported.

Type	Value	Notes
Seconds	0 - 65,535	
Minutes	0 - 65,535	
Hours	0 - 24	
Days	0 - 365	
Months	0 - 12	
Years	0 - 100	
Infinite	-	Retains the files forever.

Type	Value	Notes
		Available for Default retention, Maximum retention, and Minimum retention.
Unspecified*	-	Retains the files until you set a retention period. Available for Default retention only.

*When you transition files to WORM with an unspecified retention period, they are given the minimum retention period that is configured for the SnapLock volume. When you transition the WORM-protected files to an absolute retention time, the new retention period must be greater than the minimum period that you set on the files previously.

Expired retention period

After a WORM file's retention period expires, you can delete the file or set a new retention period to turn WORM protection back on. WORM files aren't automatically deleted after their retention period expires. You still can't modify the content of a WORM file, even after its retention period has expired.

Setting the retention period of a SnapLock volume

You can set the retention period of a SnapLock volume with the Amazon FSx console, the Amazon CLI, the Amazon FSx API, and the ONTAP CLI and REST API.

To set the retention period with the Amazon FSx API, use the [SnaplockRetentionPeriod](#) configuration.

The following procedure explains how to set the retention period on the Amazon FSx console.

To set the retention period of a SnapLock volume on the Amazon FSx console

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. Follow the procedure for creating a new volume in [Creating volumes](#).

3. In the **Advanced** section, for **SnapLock Configuration**, choose **Enabled**.

Select the check box to acknowledge the warning about enabling SnapLock on the volume.

4. For **Retention period**, enter values for **Default retention**, **Minimum retention**, and **Maximum retention**. Then choose a corresponding **Unit** for each.
5. Follow the rest of the procedure for creating a new volume in [Creating volumes](#).
6. Choose **Confirm** to create the volume.

Committing files to WORM state

This section discusses how you can transition your files to a write once, read many (WORM) state. It also discusses volume-append mode, which is a way to write data incrementally to WORM-protected files.

Autocommit

You can use autocommit to transition files to WORM if they haven't been modified for a period of time that you specify. You can turn on autocommit with the Amazon FSx console, the Amazon CLI, the Amazon FSx API, and the ONTAP CLI and REST API.

You can specify an autocommit period between five minutes and 10 years. The following table lists the specific ranges that are supported.

Unit	Value
Minutes	5 - 65,535
Hours	1 - 65,535
Days	1 - 3,650
Months	1 - 120
Years	1 - 10

To turn on autocommit with the Amazon FSx API, use `AutocommitPeriod` in the [CreateSnaplockConfiguration](#).

The following procedure explains how to turn on autocommit on the Amazon FSx console.

To turn on autocommit on the Amazon FSx console

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. Follow the procedure for creating a new volume in [Creating volumes](#).
3. In the **Advanced** section, for **SnapLock Configuration**, choose **Enabled**.

Select the check box to acknowledge the warning about enabling SnapLock on the volume.

4. For **Autocommit**, choose **Enabled**.
5. For **Autocommit period**, enter a value and choose a corresponding **Autocommit unit**.

You can specify a value between 5 minutes and 10 years.

6. Follow the rest of the procedure for creating a new volume in [Creating volumes](#).
7. Choose **Confirm** to create the volume.

Volume-append mode

You can't modify existing data in a WORM-protected file. However, SnapLock 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. You can turn volume-append mode on or off with the Amazon FSx console, the Amazon CLI, the Amazon FSx API, and the ONTAP CLI and REST API.

Requirements for updating volume-append mode

- The SnapLock volume must be unmounted.
- The SnapLock volume must be empty of snapshot copies and user data.

To turn on volume-append mode with the Amazon FSx API, use `VolumeAppendModeEnabled` in the [CreateSnaplockConfiguration](#).

The following procedure explains how to turn on volume-append mode on the Amazon FSx console.

To turn on volume-append mode on the Amazon FSx console

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. Follow the procedure for creating a new volume in [Creating volumes](#).
3. In the **Advanced** section, for **SnapLock Configuration**, choose **Enabled**.

Select the check box to acknowledge the warning about enabling SnapLock on the volume.

4. For **Volume append mode**, choose **Enabled**.
5. Follow the rest of the procedure for creating a new volume in [Creating volumes](#).
6. Choose **Confirm** to create the volume.

Event-based retention (EBR)

You can use event-based retention (EBR) to create custom policies with associated retention periods. For example, you can transition all files in a specified path to WORM and set the retention period for one year with the `snaplock event-retention policy create` and `snaplock event-retention apply` commands. When you use EBR, you must specify a volume, directory, or file. The retention period that you select when you create the EBR policy is applied to all files in the specified path.

EBR is supported by the ONTAP CLI and REST API.

Note

ONTAP doesn't support EBR with FlexGroup volumes.

The following procedures explain how to create, apply, modify, and delete an EBR policy. You must be a SnapLock administrator (have the `vsadmin-snaplock` role) to complete these tasks in the ONTAP CLI. For more information, see [SnapLock administrator](#).

To create an EBR policy in the ONTAP CLI

Run the following command. Replace *p1* and *"10 years"* with your own information.

```
vs1::> snaplock event-retention policy create -name p1 -retention-period "10 years"
```

To apply an EBR policy in the ONTAP CLI

Run the following command. Replace *p1* and *slc* with your own information. You can add a path after the forward slash (/) if you want to specify a particular path for the EBR policy. Otherwise, this command applies the EBR policy to all files on the volume.

```
vs1::> snaplock event-retention apply -policy-name p1 -volume slc -path /
```

To modify an EBR policy in the ONTAP CLI

Run the following command. Replace *p1* and "*5 years*" with your own information.

```
vs1::> snaplock event-retention policy modify -name p1 -retention-period "5 years"
```

To delete an EBR policy in the ONTAP CLI

Run the following command. Replace *p1* with your own information.

```
vs1::> snaplock event-retention policy delete -name p1
```

Related commands in the *NetApp Documentation Center*:

- [snaplock event-retention abort](#)
- [snaplock event-retention show-vservers](#)
- [snaplock event-retention show](#)
- [snaplock event-retention policy show](#)

Legal Hold

You can retain WORM files for an indefinite period of time using Legal Hold. Legal Hold is generally used for litigation purposes. A WORM file that's subject to a Legal Hold can't be deleted until the Legal Hold is lifted.

Legal Hold is supported by the ONTAP CLI and REST API.

Note

ONTAP doesn't support Legal Hold with FlexGroup volumes.

The following procedures explain how to start and end a Legal Hold. You must be a SnapLock administrator (have the `vsadmin-snaplock` role) to complete these tasks in the ONTAP CLI. For more information, see [SnapLock administrator](#).

To start a Legal Hold on a file in a SnapLock Compliance volume with the ONTAP CLI

Run the following command. Replace *litigation1*, *slc_vol1*, and *file1* with your own information.

```
vs1::> snaplock legal-hold begin -litigation-name litigation1 -volume slc_vol1 -  
path /file1
```

To start a Legal Hold on all files in a SnapLock Compliance volume with the ONTAP CLI

Run the following command. Replace *litigation1* and *slc_vol1* with your own information.

```
vs1::> snaplock legal-hold begin -litigation-name litigation1 -volume slc_vol1 -path /
```

To end a Legal Hold on a file in a SnapLock Compliance volume with the ONTAP CLI

Run the following command. Replace *litigation1*, *slc_vol1*, and *file1* with your own information.

```
vs1::> snaplock legal-hold end -litigation-name litigation1 -volume slc_vol1 -  
path /file1
```

To end a Legal Hold on all files in a SnapLock Compliance volume with the ONTAP CLI

Run the following command. Replace *litigation1* and *slc_vol1* with your own information.

```
vs1::> snaplock legal-hold end -litigation-name litigation1 -volume slc_vol1 -path /
```

Note

We recommend that you monitor the `-operation-status` with the `snaplock legal-hold show` command when issuing a Legal Hold to make sure that it doesn't fail.

Related commands in the *NetApp Documentation Center*:

- [snaplock legal-hold abort](#)
- [snaplock legal-hold dump-files](#)
- [snaplock legal-hold dump-litigations](#)
- [snaplock legal-hold show](#)

Backing up SnapLock volumes

You can back up SnapLock volumes for additional data protection. When you restore a SnapLock volume, the volume's original settings—such as the default retention, minimum retention, and maximum retention—are preserved. Write once, read many (WORM) settings and Legal Hold are also preserved.

Note

You can't back up a SnapLock FlexGroup volume.

You can restore a SnapLock volume's backup as a SnapLock or a non-SnapLock volume. However, you can't restore a non-SnapLock volume's backup as a SnapLock volume.

For more information about backups, see [Working with backups](#).

Deleting SnapLock volumes


You can delete a SnapLock Compliance volume if the retention periods of all the write once, read many (WORM) files on it are expired.

Note

When you close an Amazon Web Services account that contains SnapLock Enterprise or Compliance volumes, Amazon and FSx for ONTAP suspend your account for 90 days with your data intact. If you don't reopen your account during those 90 days, Amazon deletes your data including data in SnapLock volumes regardless of your retention settings.

You can delete a SnapLock Enterprise volume at any time if you have the appropriate permissions. You must be an Amazon FSx administrator. Additionally, whether

you're using the Amazon FSx console or Amazon FSx API, you must have the IAM `fsx:BypassSnapLockEnterpriseRetention` IAM permission to delete a SnapLock Enterprise volume that contains WORM data with an active retention policy.

 **Warning**

The minimum retention period for a SnapLock audit log volume is six months. Until this retention period expires you can't delete the SnapLock audit log volume, the storage virtual machine (SVM), or the file system that's associated with the SVM—even if the volume was created in SnapLock Enterprise mode. For more information, see [SnapLock audit log volumes](#).

To delete a SnapLock Enterprise volume on the Amazon FSx console

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. In the left navigation pane, choose **Volumes**.
3. Choose the volume that you want to delete.
4. From **Actions**, choose **Delete volume**.
5. For **Bypass SnapLock Enterprise Retention**, choose **Yes**.
6. In the confirmation dialog box, choose one of the following options for **Create final backup**:
 - Choose **Yes** to make a final backup of the volume. The name of the final backup is displayed.
 - Choose **No** if you don't want a final backup of the volume. You are asked to acknowledge that once the volume is deleted, automatic backups are no longer available.
7. Confirm the volume deletion by entering **delete** in the **Confirm delete** field.
8. Choose **Delete volume(s)**.

Working with Microsoft Active Directory in FSx for ONTAP

Amazon FSx works with Microsoft Active Directory to integrate with your existing environments. Active Directory is the Microsoft directory service that's used to store information about objects on the network, and to help administrators and users to find and use this information. These objects typically include shared resources, such as file servers and network user and computer accounts.

You can optionally join your FSx for ONTAP storage virtual machines (SVMs) to your Active Directory domain to provide user authentication and file- and folder-level access control. Server message block (SMB) clients can then use their existing user identities in Active Directory to authenticate themselves and access SVM volumes. Your users can use their existing identities to control access to individual files and folders. In addition, you can migrate your existing files and folders and their security access control list (ACL) configurations to Amazon FSx without any modifications.

When you join Amazon FSx for NetApp ONTAP to an Active Directory, you join the file system's SVMs to the Active Directory independently. This means that you can have a file system with some SVMs that are joined to an Active Directory, and other SVMs that are not.

After an SVM is joined to an Active Directory, you can update the following Active Directory configuration properties:

- DNS server IP addresses
- Self-managed Active Directory service account username and password

Topics

- [Prerequisites for joining an SVM to a self-managed Microsoft AD](#)
- [Best practices for working with Active Directory](#)
- [Joining SVMs to a Microsoft Active Directory](#)
- [Managing SVM Active Directory configurations](#)

Prerequisites for joining an SVM to a self-managed Microsoft AD

Before you join an FSx for ONTAP SVM to a self-managed Microsoft AD domain, make sure that your Active Directory and network meet the requirements described in the following sections.

Topics

- [On-premises Active Directory requirements](#)
- [Network configuration requirements](#)
- [Active Directory service account requirements](#)

On-premises Active Directory requirements

Make sure that you already have an on-premises or other self-managed Microsoft AD that you can join the SVM to. This Active Directory should have the following configuration:

- The Active Directory domain controller domain functional level is at Windows Server 2000 or higher.
- The Active Directory uses a domain name that's not in the Single Label Domain (SLD) format. Amazon FSx doesn't support SLD domains.
- If you have Active Directory sites defined, make sure that the subnets in the VPC that's associated with your FSx for ONTAP file system are defined in the same Active Directory sites, and that no conflicts exist between your VPC subnets and the subnets on your Active Directory sites.

Note

If you are using Amazon Directory Service, FSx for ONTAP doesn't support joining SVMs to the Simple Active Directory.

Network configuration requirements

Make sure that you have the following network configurations in place and associated information available to you.


⚠ Important

For an SVM to join Active Directory, you need to ensure that the ports documented in this topic allow traffic between all Active Directory Domain Controllers and both iSCSI IP addresses (iscsi_1 and iscsi_2 logical interfaces (LIFs)) on the SVM.

- The DNS server and Active Directory domain controller IP addresses.
- Connectivity between the Amazon VPC where you're creating the file system and your self-managed Active Directory using [Amazon Direct Connect](#), [Amazon VPN](#), or [Amazon Transit Gateway](#).
- The security group and the VPC Network ACLs for the subnets on which you're creating the file system must allow traffic on the ports and in the directions shown in the following diagram.

FSx for ONTAP File Server port requirements

Configure VPC security groups that you've associated with your Amazon FSx file system, along with any VPC Network ACLs and ONTAP firewalls to allow network traffic on the following ports:



The role of each port is described in the following table.

Protocol	Ports	Role
TCP/UDP	53	Domain Name System (DNS)
TCP/UDP	88	Kerberos authentication
TCP/UDP	389	Lightweight Directory Access Protocol (LDAP)
TCP	445	Directory Services SMB file sharing
TCP/UDP	464	Change/Set password
TCP	636	Lightweight Directory Access Protocol over TLS/SSL (LDAPS)

- These traffic rules should also be mirrored on the firewalls that apply to each of the Active Directory domain controllers, DNS servers, FSx clients, and FSx administrators.

Important

While Amazon VPC security groups require ports to be opened only in the direction that network traffic is initiated, most Windows firewalls and VPC network ACLs require ports to be open in both directions.

Active Directory service account requirements

Make sure that you have a service account in your self-managed Microsoft AD that has delegated permissions to join computers to the domain. A *service account* is a user account in your self-managed Active Directory that has been delegated certain tasks.

At a minimum, the service account must be delegated the following permissions in the OU to which you're joining the SVM:

- Ability to reset passwords
- Ability to restrict accounts from reading and writing data
- Ability to set the `msDS-SupportedEncryptionTypes` property on computer objects
- Validated ability to write to the DNS hostname

- Validated ability to write to the service principal name
- Ability to create and delete computer objects
- Validated ability to read and write Account Restrictions

These represent the minimum set of permissions that are required to join computer objects to your Active Directory. For more information, see the Windows Server documentation topic [Error: Access is denied when non-administrator users who have been delegated control try to join computers to a domain controller](#).

To learn more about creating a service account with the correct permissions, see [Delegating permissions to your Amazon FSx service account](#).

⚠ Important

Amazon FSx requires a valid service account throughout the lifetime of your Amazon FSx file system. Amazon FSx must be able to fully manage the file system and perform tasks that require it to unjoin and rejoin resources to your Active Directory domain. These tasks include replacing a failed file system or SVM, or patching NetApp ONTAP software. Keep your Active Directory configuration information up to date with Amazon FSx, including the service account credentials. To learn more, see [Keeping your Active Directory configuration updated with Amazon FSx](#).

If this is your first time using Amazon and FSx for ONTAP, make sure that you complete the initial setup steps before starting your Active Directory integration. For more information, see [Setting up FSx for ONTAP](#).

⚠ Important

Don't move computer objects that Amazon FSx creates in the OU after your SVMs are created, or delete your Active Directory while your SVM is joined to it. Doing so will cause your SVMs to become misconfigured.

Best practices for working with Active Directory

Here are some suggestions and guidelines that you should consider when joining Amazon FSx for NetApp ONTAP SVMs to your self-managed Microsoft Active Directory. Note that these are recommended as best practices, but not required.

Delegating permissions to your Amazon FSx service account

Make sure to configure the service account that you provide to Amazon FSx with the minimum permissions required. In addition, separate the Organizational Unit (OU) from other domain controller concerns.

To join Amazon FSx SVMs to your domain, make sure that the service account has delegated permissions. Members of the **Domain Admins** group have sufficient permissions to perform this task. However, as a best practice, use a service account that only has the minimum permissions necessary to do this. The following procedure demonstrates how to delegate only the permissions necessary to join FSx for ONTAP SVMs to your domain.

Perform this procedure on a machine that's joined to your directory and has the Active Directory User and Computers MMC snap-in installed.

To create a service account for your Microsoft Active Directory domain

1. Make sure that you're logged in as a domain administrator for your Microsoft Active Directory domain.
2. Open the **Active Directory User and Computers** MMC snap-in.
3. In the task pane, expand the domain node.
4. Locate and open the context (right-click) menu for the OU that you want to modify, and then choose **Delegate Control**.
5. On the **Delegation of Control Wizard** page, choose **Next**.
6. Choose **Add** to add a specific user or a specific group for **Selected users and groups**, and then choose **Next**.
7. On the **Tasks to Delegate** page, choose **Create a custom task to delegate**, and then choose **Next**.
8. Choose **Only the following objects in the folder**, and then choose **Computer objects**.
9. Choose **Create selected objects in this folder** and **Delete selected objects in this folder**. Then choose **Next**.

10. Under **Show these permissions**, ensure that **General** and **Property-specific** are selected.
11. For **Permissions**, choose the following:
 - **Reset Password**
 - **Read and write Account Restrictions**
 - **Validated write to DNS host name**
 - **Validated write to service principal name**
 - **Write msDS-SupportedEncryptionTypes**
12. Choose **Next**, and then choose **Finish**.
13. Close the **Active Directory User and Computers** MMC snap-in.

Important

Don't move computer objects that Amazon FSx creates in the OU after your SVMs are created. Doing so will cause your SVMs to become misconfigured.

Keeping your Active Directory configuration updated with Amazon FSx

For uninterrupted availability of your Amazon FSx SVMs, update an SVM's self-managed Active Directory (AD) configuration when you change your self-managed AD setup.

For example, suppose that your AD uses a time-based password reset policy. In this case, as soon as the password is reset, make sure to update the service account password with Amazon FSx. To do this, use the Amazon FSx console, Amazon FSx API, or Amazon CLI. Similarly, if the DNS server IP addresses change for your Active Directory domain, as soon as the change occurs update the DNS server IP addresses with Amazon FSx.

If there's an issue with the updated self-managed AD configuration, the SVM state switches to **Misconfigured**. This state shows an error message and a recommended action beside the SVM description in the console, API, and CLI. If an issue with your SVM's AD configuration occurs, be sure to take the recommended corrective action for the configuration properties. If the issue is resolved, verify that your SVM's state changes to **Created**.

For more information, see [Updating an existing SVM Active Directory configuration using the Amazon Web Services Management Console, Amazon CLI, and API](#) and [Modify an Active Directory configuration using the ONTAP CLI](#).

Using security groups to limit traffic within your VPC

To limit network traffic in your virtual private cloud (VPC), you can implement the principle of least privilege in your VPC. In other words, you can limit permissions to the minimum ones necessary. To do this, use security group rules. To learn more, see [Amazon VPC security groups](#).

Creating outbound security group rules for your file system's network interface

For greater security, consider configuring a security group with outbound traffic rules. These rules should allow outbound traffic only to your self-managed AD domains controllers or within the subnet or security group. Apply this security group to the VPC associated with your Amazon FSx file system's elastic network interface. To learn more, see [File System Access Control with Amazon VPC](#).

Joining SVMs to a Microsoft Active Directory

Your organization might manage identities and devices using an Active Directory, whether on-premises or in the cloud. With FSx for ONTAP, you can join your SVMs directly to your existing Active Directory domain in the following ways:

- Joining new SVMs to an Active Directory at creation:
 - Using the **Standard create** option in Amazon FSx console to create a new FSx for ONTAP file system, you can join the default SVM to a self-managed Active Directory. For more information, see [To create a file system \(console\)](#).
 - Using the Amazon FSx console, Amazon CLI, or Amazon FSx API to create a new SVM on an existing FSx for ONTAP file system. For more information, see [Creating a storage virtual machine](#).
- Joining existing SVMs to an Active Directory:
 - Using the Amazon Web Services Management Console, Amazon CLI, and API to join an SVM to an Active Directory, and to reattempt joining an SVM to an Active Directory if the initial attempt to join failed. You can also update some Active Directory configuration properties for SVMs that are already joined to an Active Directory. For more information, see [Managing SVM Active Directory configurations](#).
 - Using the NetApp ONTAP CLI or REST API to join, reattempt joining, and unjoining SVM Active Directory configurations. For more information, see [Managing your SVM Active Directory configuration using the NetApp CLI](#).

Important

- Amazon FSx only registers DNS records for an SVM if you use Microsoft DNS as the default DNS service. If you use a third-party DNS, you must set up DNS entries manually for your Amazon FSx SVMs after you create them.
- If you use Amazon Managed Microsoft AD, you must specify a group such as Amazon Delegated FSx Administrators, Amazon Delegated Administrators, or a custom group with delegated permissions to the OU.

When you join an FSx for ONTAP SVM directly to a self-managed Active Directory, the SVM resides in the same Active Directory forest (the top-most logical container in an Active Directory configuration that contains domains, users, and computers) and in the same Active Directory domain as your users and existing resources, including existing file servers.

Information needed when joining an SVM to an Active Directory

You have to provide the following information about your Active Directory when joining an SVM to an Active Directory, regardless of the API operation you choose:

- The NetBIOS name of the Active Directory computer object to create for your SVM. This is the name of the SVM in Active Directory, which must be unique within your Active Directory. Don't use the NetBIOS name of the home domain. The NetBIOS name can't exceed 15 characters.
- The fully qualified domain name (FQDN) of your Active Directory. The FQDN can't exceed 255 characters.

Note

The FQDN can't be in the Single Label Domain (SLD) format. Amazon FSx doesn't support SLD domains.

- Up to three IP addresses of the DNS servers or domain hosts for your domain.

The DNS server IP addresses and Active Directory domain controller IP addresses can be in any IP address range, except:

- IP addresses that conflict with Amazon Web Services–owned IP addresses in that Amazon Web Services Region. For a list of Amazon IP addresses by Region, see the [Amazon IP address ranges](#).
- IP addresses in the following CIDR block range: 198.19.0.0/16
- Username and password for a service account on your Active Directory domain for Amazon FSx to use when joining the SVM to the Active Directory domain. For more information about service account requirements, see [Active Directory service account requirements](#).
- (Optional) The Organizational Unit (OU) in the domain that you join the SVM to.

Note

If you join your SVM to an Amazon Directory Service Active Directory, you must provide an OU that's within the default OU that Amazon Directory Service creates for the directory objects that are related to Amazon. This is because the Amazon Directory Service doesn't provide access to your Active Directory's default Computers OU. For example, if your Active Directory domain is `example.com`, you can specify the following OU: `OU=Computers,OU=example,DC=example,DC=com`.

- (Optional) The domain group that you are delegating authority to for performing administrative actions on your file system. For example, this domain group might manage Windows SMB file shares, take ownership of files and folders, and so on. If you don't specify this group, Amazon FSx delegates this authority to the Domain Admins group in your Active Directory domain by default.

Managing SVM Active Directory configurations

This section describes how to use the Amazon Web Services Management Console, Amazon CLI, FSx API, and the ONTAP CLI to do the following:

- Joining an existing SVM to an Active Directory
- Modifying an existing SVM Active Directory configuration
- Removing SVMs from an Active Directory

To remove an SVM from an Active Directory, you must use the NetApp ONTAP CLI.

Topics

- [Joining an SVM to an Active Directory using the Amazon Web Services Management Console, Amazon CLI and API](#)
- [Updating an existing SVM Active Directory configuration using the Amazon Web Services Management Console, Amazon CLI, and API](#)
- [Managing your SVM Active Directory configuration using the NetApp CLI](#)

Joining an SVM to an Active Directory using the Amazon Web Services Management Console, Amazon CLI and API

Use the following procedure to join an existing SVM to an Active Directory. In this procedure, the SVM is *not* already joined to an Active Directory.

To join an SVM to an Active Directory (Amazon Web Services Management Console)

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. Choose the SVM that you want to join to an Active Directory:
 - In the left navigation pane, choose **File systems**, and then choose the ONTAP file system with the SVM that you want to update.
 - Choose the **Storage virtual machines** tab.

–Or–

 - To display a list of all of the available SVMs, in the left navigation pane, expand **ONTAP** and choose **Storage virtual machines**. A list of all SVMs in your account in the Amazon Web Services Region is displayed.

Select the SVM that you want to join to an Active Directory from the list.
3. On the upper right of the SVM **Summary** panel, choose **Actions > Join/Update Active Directory**. The **Join SVM to an Active Directory** window appears.
4. Enter the following information for the Active Directory that you are joining the SVM to:
 - The **NetBIOS name** of the Active Directory computer object to create for your SVM. This is the name of the SVM in Active Directory, which must be unique within your Active Directory. Don't use the NetBIOS name of the home domain. The NetBIOS name can't exceed 15 characters.

- The **fully qualified domain name (FQDN)** of your Active Directory. The domain name can't exceed 255 characters.
- **DNS server IP addresses** – The IPv4 addresses of the DNS servers for your domain.
- **Service account username** – The user name of the service account in your existing Active Directory. Don't include a domain prefix or suffix. For example, for EXAMPLE\ADMIN, use only ADMIN.
- **Service account password** – The password for the service account.
- **Confirm password** – The password for the service account.
- (Optional) **Organizational Unit (OU)** – The distinguished path name of the organizational unit you want to join your SVM to.
- **Delegated file system administrators group** – The name of the group in your Active Directory that can administer your file system.

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

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

5. Choose **Join Active Directory** to join the SVM to the Active Directory using the configuration you provided.

To join an SVM to an Active Directory (Amazon CLI)

- To join an FSx for ONTAP SVM to an Active Directory, use the [update-storage-virtual-machine](#) CLI command (or the equivalent [UpdateStorageVirtualMachine](#) API operation), as shown in the following example.

```
aws fsx update-storage-virtual-machine \
  --storage-virtual-machine-id svm-abcdef0123456789a\
  --active-directory-configuration
  SelfManagedActiveDirectoryConfiguration='{DomainName="corp.example.com", \
    OrganizationalUnitDistinguishedName="OU=FileSystems,DC=corp,DC=example,DC=com",
  \
    FileSystemAdministratorsGroup="FSxAdmins",UserName="FSxService",\
    Password="password", \
    DnsIps=["10.0.1.18"]}',NetBiosName=amznfsx12345
```


After successfully creating the storage virtual machine, Amazon FSx returns its description in JSON format, as shown in the following example.

```
{
  "StorageVirtualMachine": {
    "ActiveDirectoryConfiguration": {
      "NetBiosName": "amznfsx12345",
      "SelfManagedActiveDirectoryConfiguration": {
        "UserName": "Admin",
        "DnsIps": [
          "10.0.1.3",
          "10.0.91.97"
        ],
        "OrganizationalUnitDistinguishedName": "OU=Computers,OU=customer-
ad,DC=customer-ad,DC=example,DC=com",
        "DomainName": "customer-ad.example.com"
      }
    }
  },
  "CreationTime": 1625066825.306,
  "Endpoints": {
    "Management": {
      "DnsName": "svm-abcdef0123456789a.fs-0123456789abcdef0.fsx.us-
east-1.amazonaws.com",
      "IpAddresses": ["198.19.0.4"]
    },
    "Nfs": {
      "DnsName": "svm-abcdef0123456789a.fs-0123456789abcdef0.fsx.us-
east-1.amazonaws.com",
      "IpAddresses": ["198.19.0.4"]
    },
    "Smb": {
      "DnsName": "amznfsx12345",
      "IpAddresses": ["198.19.0.4"]
    },
    "SmbWindowsInterVpc": {
      "IpAddresses": ["198.19.0.5", "198.19.0.6"]
    },
    "Iscsi": {
      "DnsName": "iscsi.svm-abcdef0123456789a.fs-0123456789abcdef0.fsx.us-
east-1.amazonaws.com",
      "IpAddresses": ["198.19.0.7", "198.19.0.8"]
    }
  }
}
```

```
    },
    "FileSystemId": "fs-0123456789abcdef0",
    "Lifecycle": "CREATED",
    "Name": "vol1",
    "ResourceARN": "arn:aws:fsx:us-east-1:123456789012:storage-virtual-machine/
fs-0123456789abcdef0/svm-abcdef0123456789a",
    "StorageVirtualMachineId": "svm-abcdef0123456789a",
    "Subtype": "default",
    "Tags": [],

  }
}
```

Updating an existing SVM Active Directory configuration using the Amazon Web Services Management Console, Amazon CLI, and API

Use the following procedure to update the Active Directory configuration of an SVM that is already joined to an Active Directory.

To update an SVM Active Directory configuration (Amazon Web Services Management Console)

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. Choose the SVM to update as follows:
 - In the left navigation pane, choose **File systems**, and then choose the ONTAP file system with the SVM you want to update.
 - Choose the **Storage virtual machines** tab.

–Or–

- To display a list of all of the SVMs available, in the left navigation pane, expand **ONTAP** and choose **Storage virtual machines**.

Select the SVM that you want to update from the list.

3. On the SVM **Summary** panel, choose **Actions** > **Join/Update Active Directory**. The **Update SVM Active Directory configuration** window appears.
4. You can update the following Active Directory configuration properties in this window.
 - **DNS server IP addresses** – The IPv4 addresses of the DNS servers for your domain.

- **Service account username** – The username of the service account in your existing Active Directory. Don't include a domain prefix or suffix. For EXAMPLE\ADMIN, use ADMIN.
 - **Service account password** – The password for the Active Directory service account.
5. After you have entered your updates, choose **Update Active Directory** to make the changes.

Use the following procedure to update the Active Directory configuration of an SVM that is already joined to an Active Directory.

To update an SVM Active Directory configuration (Amazon CLI)

- To update an SVM's Active Directory configuration with the Amazon CLI or API, use the [update-storage-virtual-machine](#) CLI command (or the equivalent [UpdateStorageVirtualMachine](#) API operation), as shown in the following example.

```
aws fsx update-storage-virtual-machine \  
  --storage-virtual-machine-id svm-abcdef0123456789a\  
  --active-directory-configuration \  
  SelfManagedActiveDirectoryConfiguration='{UserName="FSxService",\  
  Password="password", \  
  DnsIps=["10.0.1.18"]}'
```

Managing your SVM Active Directory configuration using the NetApp CLI

You can use the NetApp ONTAP CLI to join and unjoin your SVM to an Active Directory, and to modify an existing SVM Active Directory configuration.

Joining an SVM to an Active Directory using the ONTAP CLI

You can join existing SVMs to an Active Directory using the ONTAP CLI, as described in the following procedure. You can do this even if your SVM is already joined to an Active Directory.

1. To access the NetApp ONTAP CLI, establish an SSH session on the management port of the Amazon FSx for NetApp ONTAP file system by running the following command. Replace *management_endpoint_ip* with the IP address of the file system's management port.

```
[~]$ ssh fsxadmin@management_endpoint_ip
```

For more information, see [Managing file systems with the ONTAP CLI](#).

2. Create a DNS entry for your Active Directory by providing the full directory DNS name (corp.example.com) and at least one DNS server IP address.

```
::>vserver services name-service dns create -vserver svm_name -
domains corp.example.com -name-servers dns_ip_1, dns_ip_2
```

To verify the connection to your DNS servers, run the following command. Replace *svm_name* with your own information.

```
FsxId0ae30e5b7f1a50b6a:>vserver services name-service dns check -vserver svm_name
```

Vserver	Name Server	Name Server Status	Status Details
<i>svm_name</i>	172.31.14.245	up	Response time (msec): 0
<i>svm_name</i>	172.31.25.207	up	Response time (msec): 1

2 entries were displayed.

3. To join your SVM to your Active Directory, run the following command. Note that you will must specify a *computer_name* that doesn't already exist in your Active Directory and provide the directory DNS name for *-domain*. For *-OU*, enter the OUs that you want the SVM to join, as well as the full DNS name in DC format.

```
::>vserver cifs create -vserver svm_name -cifs-server computer_name -
domain corp.example.com -OU OU=Computers,OU=example,DC=corp,DC=example,DC=com
```

To verify the status of your Active Directory connection, run the following command:

```
::>vserver cifs check -vserver svm_name
```

```

Vserver : svm_name
      Cifs NetBIOS Name : svm_netBIOS_name
      Cifs Status : Running
      Site : Default-First-Site-Name
Node Name      DC Server Name  DC Server IP   Status   Status Details
-----
FsxId0ae30e5b7f1a50b6a-01
      corp.example.com
```

```

                                172.31.14.245   up           Response time (msec): 5
FsxId0ae30e5b7f1a50b6a-02
                                corp.example.com
                                172.31.14.245   up           Response time (msec): 20
2 entries were displayed.
```

4. If you can't access shares after this join, determine whether the account you're using to access the share has permissions. For example, if you're using the default Admin account (a delegated administrator) with an Amazon managed Active Directory, you will must run the following command in ONTAP. The `netbios_domain` corresponds with your Active Directory's domain name (for `corp.example.com`, the `netbios_domain` used here is `example`).

```
FsxId0123456789a::>vserver cifs users-and-groups local-group add-members -vserver
svm_name -group-name BUILTIN\Administrators -member-names netbios_domain\admin
```

Modify an Active Directory configuration using the ONTAP CLI

You can use the ONTAP CLI to modify an existing Active Directory configuration.

1. To access the NetApp ONTAP CLI, establish an SSH session on the management port of the Amazon FSx for NetApp ONTAP file system by running the following command. Replace *management_endpoint_ip* with the IP address of the file system's management port.

```
[~]$ ssh fsxadmin@management_endpoint_ip
```

For more information, see [Managing file systems with the ONTAP CLI](#).

2. Run the following command to temporarily bring down the SVM's CIFS server:

```
FsxId0123456789a::>vserver cifs modify -vserver svm_name -status-admin down
```

3. If you need to modify the DNS entries of your Active Directory, run the following command:

```
::>vserver services name-service dns modify -vserver svm_name -
domains corp.example.com -name-servers dns_ip_1,dns_ip_2
```

You can validate the connection status to your Active Directory's DNS servers using the `vserver services name-service dns check -vserver svm_name` command.

```
::>vserver services name-service dns check -vserver svm_name
```

Vserver	Name Server	Status	Status Details
svmciaad	dns_ip_1	up	Response time (msec): 1
svmciaad	dns_ip_2	up	Response time (msec): 1

2 entries were displayed.

4. If you need to modify the Active Directory configuration itself, you can change existing fields by using the following command, replacing:

- *computer_name*, if you want to modify the NetBIOS (machine account) name of the SVM.
- *domain_name*, if you want to modify the name of the domain. This should correspond with the DNS domain entry noted in Step 3 of this section (`corp.example.com`).
- `organizational_unit`, if you want to modify the OU (`OU=Computers,OU=example,DC=corp,DC=example,DC=com`).

You will need to reenter the Active Directory credentials that you used to join this device to the Active Directory.

```
::>vserver cifs modify -vserver svm_name -cifs-server computer_name -  
domain domain_name -OU organizational_unit
```

You can verify the connection status of your Active Directory connection using the `vserver cifs check -vserver svm_name` command.

5. When you finish modifying your Active Directory and DNS configuration, bring the CIFS server back up by running the following command:

```
::>vserver cifs modify -vserver svm_name -status-admin up
```

Unjoin an Active Directory from your SVM using the NetApp ONTAP CLI

The NetApp ONTAP CLI can also be used to unjoin your SVM from an Active Directory by following the steps below:

1. To access the NetApp ONTAP CLI, establish an SSH session on the management port of the Amazon FSx for NetApp ONTAP file system by running the following command. Replace *management_endpoint_ip* with the IP address of the file system's management port.

```
[~]$ ssh fsxadmin@management_endpoint_ip
```

For more information, see [Managing file systems with the ONTAP CLI](#).

2. Delete the CIFS server that unjoined your device from the Active Directory by running the following command. For ONTAP to delete the machine account for your SVM, provide the credentials that you originally used to join the SVM to the Active Directory.

```
FsxId0123456789a::>vserver cifs modify -vserver svm_name -status-admin down
```

3. If you need to modify the DNS entries of your Active Directory, run the following command:

```
FsxId0123456789a::vserver cifs delete -vserver svm_name
```

In order to delete an Active Directory machine account for the CIFS server, you must supply the name and password of a Windows account with sufficient privileges to remove computers from the "CORP.ADEXAMPLE.COM" domain.

Enter the user name: *user_name*

Enter the password:

Warning: There are one or more shares associated with this CIFS server
Do you really want to delete this CIFS server and all its shares? {y|n}: *y*

4. Delete the DNS servers for your Active Directory by running the following command:

```
::vserver services name-service dns delete -vserver svm_name
```

If you see a warning like the following—indicating that dns should be removed as an ns-switch—and you don't plan to rejoin this device to an Active Directory, you can remove the ns-switch entries.

```
Warning: "DNS" is present as one of the sources in one or more ns-switch databases  
but no valid DNS configuration was found for Vserver  
"svm_name". Remove "DNS" from ns-switch using the "vserver services name-  
service ns-switch" command. Configuring "DNS" as a source
```

in the ns-switch setting when there is no valid configuration can cause protocol access issues.

5. (Optional) Remove the ns-switch entries for dns by running the following command. Verify the source order, then remove the dns entry for the hosts database by modifying the sources so that they contain only the other sources listed. In this example, the only other source is files.

```
::>vserver services name-service ns-switch show -vserver svm_name -database hosts
```

```
          Vserver: svm_name
Name Service Switch Database: hosts
  Name Service Source Order: files, dns
```

```
::>vserver services name-service ns-switch modify -vserver svm_name -database hosts
-sources files
```

6. (Optional) Remove the dns entry by modifying the sources for the database host to include only files.

```
::>vserver services name-service ns-switch modify -vserver svm_name -database hosts
-sources files
```


Amazon FSx for NetApp ONTAP performance

Following is an overview of Amazon FSx for NetApp ONTAP file system performance, with a discussion of the available performance and throughput options and useful performance tips.

Topics

- [How performance is measured for FSx for ONTAP file systems](#)
- [Performance details](#)
- [Impact of deployment type on performance](#)
- [Impact of storage capacity on performance](#)
- [Impact of throughput capacity on performance](#)
- [Example: storage capacity and throughput capacity](#)

How performance is measured for FSx for ONTAP file systems

File system performance is measured by its latency, throughput, and I/O operations per second (IOPS).

Latency

Amazon FSx for NetApp ONTAP provides sub-millisecond file operation latencies with solid state drive (SSD) storage, and tens of milliseconds of latency for capacity pool storage. Above that, Amazon FSx has two layers of read caching on each file server—NVMe (non-volatile memory express) drives and in-memory—to provide even lower latencies when you access your most frequently-read data.

Throughput and IOPS

Each Amazon FSx file system provides up to tens of GB/s of throughput and millions of IOPS. The specific amount of throughput and IOPS that your workload can drive on your file system depends on the total throughput capacity and storage capacity configuration of your file system, along with the nature of your workload, including the size of the active working set.

SMB Multichannel and NFS nconnect support


With Amazon FSx, you can configure SMB Multichannel to provide multiple connections between ONTAP and clients in a single SMB session. SMB Multichannel uses multiple network connections between the client and server simultaneously to aggregate network bandwidth for maximal utilization. For information on using the NetApp ONTAP CLI to configure SMB Multichannel, see [Configuring SMB Multichannel for performance and redundancy](#).

NFS clients can use the nconnect mount option to have multiple TCP connections (up to 16) associated with a single NFS mount. Such an NFS client multiplexes file operations onto multiple TCP connections in a round-robin fashion and thus obtains higher throughput from the available network bandwidth. NFSv3 and NFSv4.1+ support nconnect. [Amazon EC2 instance network bandwidth](#) describes the full duplex 5 Gbps per network flow bandwidth limit. You can overcome this limit by using multiple network flows with nconnect or SMB multichannel. See your NFS client documentation to confirm whether nconnect is supported in your client version. For more information about NetApp ONTAP support for nconnect, see [ONTAP support for NFSv4.1](#).

Performance details

To understand the Amazon FSx for NetApp ONTAP performance model in detail, you can examine the architectural components of an Amazon FSx file system. Your client compute instances, whether they exist in Amazon or on-premises, access your file system through one or multiple elastic network interfaces (ENI). These network interfaces reside in the Amazon VPC that you associate with your file system. Behind each file system ENI is an NetApp ONTAP file server that is serving data over the network to the clients accessing the file system. Amazon FSx provides a fast in-memory cache and NVMe cache on each file server to enhance performance for the most frequently accessed data. Attached to each file server are the SSD disks hosting your file system data.

These components are illustrated in the following diagram.



Corresponding with these architectural components—network interface, in-memory cache, NVMe cache, and storage volumes—are the primary performance characteristics of an Amazon FSx for NetApp ONTAP file system that determine the overall throughput and IOPS performance.

- Network I/O performance: throughput/IOPS of requests between the clients and the file server (in aggregate)
- In-memory and NVMe cache size on the file server: size of active working set that can be accommodated for caching
- Disk I/O performance: throughput/IOPS of requests between the file server and the storage disks

There are two factors that determine these performance characteristics for your file system: the total amount of SSD IOPS and throughput capacity that you configure for it. The first two performance characteristics – network I/O performance and in-memory and NVMe cache size – are solely determined by throughput capacity, while the third one – disk I/O performance – is determined by a combination of throughput capacity and SSD IOPS.

File-based workloads are typically spiky, characterized by short, intense periods of high I/O with plenty of idle time between bursts. To support spiky workloads, in addition to the baseline speeds

that a file system can sustain 24/7, Amazon FSx provides the capability to burst to higher speeds for periods of time for both network I/O and disk I/O operations. Amazon FSx uses a network I/O credit mechanism to allocate throughput and IOPS based on average utilization — file systems accrue credits when their throughput and IOPS usage is below their baseline limits, and can use these credits when they perform I/O operations.

Write operations use twice as much network bandwidth as read operations. A write operation has to be replicated on the secondary file server, so a single write operation results in twice the amount of network throughput.

Impact of deployment type on performance

You can create two types of file systems with FSx for ONTAP. File systems with a single high-availability (HA) pair of file servers are called *scale-up* file systems. File systems with multiple HA pairs are called *scale-out* file systems. For more information, see [High-availability \(HA\) pairs](#).

FSx for ONTAP Multi-AZ and Single-AZ file systems provide consistent sub-millisecond file operation latencies with SSD storage and tens of milliseconds of latency with capacity pool storage. Additionally, file systems that meet the following requirements provide an NVMe read cache to reduce read latencies and increase IOPS for frequently-read data:

- Multi-AZ file systems
- Single-AZ scale-up file systems created after November 28, 2022 with at least 2 GBps of throughput capacity

The following tables show the amount of throughput capacity that file systems can scale up to depending on factors such as the number of high availability (HA) pairs and Amazon Web Services Regions availability.

Scale-up

These performance specifications apply to scale-up file systems.

Maximum throughput from SSD storage per HA pair for scale-up file systems

	US East (Ohio) Region, US East (N. Virginia) Region, US West (Oregon) Region, and Europe (Ireland)		All other Amazon Web Services Regions where FSx for ONTAP is available	
	Read throughput (MBps)	Write throughput (MBps)	Read throughput (MBps)	Write throughput (MBps)
Single-AZ	4,096*	1,000	2,048	750
Multi-AZ	4,096*	1,800	2,048	1,300

Note

* To provision 4 GBps of throughput capacity, your file system must be configured with a minimum of 5,120 GiB of SSD storage capacity and 160,000 SSD IOPS.

Scale-out

These performance specifications apply to scale-out file systems.

Maximum throughput from SSD storage per HA pair for scale-out file systems

	Read throughput (MBps)	Write throughput (MBps)
Single-AZ scale-out	6,144*	1,100*

Note

* Per HA pair (up to 12). For more information, see [High-availability \(HA\) pairs](#).

Impact of storage capacity on performance

The maximum disk throughput and IOPS levels your file system can achieve is the lower of:

- the disk performance level provided by your file servers, based on the throughput capacity you select for your file system
- the disk performance level provided by the number of SSD IOPS you provision for your file system

By default, your file system's SSD storage provides up to the following levels of disk throughput and IOPS:

- Disk throughput (MBps per TiB of storage): 768
- Disk IOPS (IOPs per TiB of storage): 3,072

Impact of throughput capacity on performance

Every Amazon FSx file system has a throughput capacity that you configure when the file system is created. Your file system's throughput capacity determines the level of network I/O performance, or the speed at which each of the file servers that are hosting your file system can serve file data over the network to clients accessing it. Higher levels of throughput capacity come with more memory and non-volatile memory express (NVMe) storage for caching data on each file server, and higher levels of disk I/O performance supported by each file server.

You can optionally provision a higher level of SSD IOPS when creating your file system. The maximum level of SSD IOPS that your file system can achieve is also dictated by your file system's throughput capacity, even when provisioning additional SSD IOPS.

The following tables show the full set of specifications for throughput capacity, along with baseline and burst levels, and amount of memory for caching on the file server in the corresponding Amazon Web Services Regions.

Single-AZ (scale-up)

These performance specifications apply to Single-AZ scale-up file systems created after November 28, 2022 in the specified Amazon Web Services Regions.

Performance specifications for file systems in the following Amazon Web Services Regions: US East (N. Virginia), US East (Ohio), US West (Oregon), and Europe (Ireland)

FSx throughput capacity (MBps)	Network throughput capacity (MBps)		Network In- IOPS memory caching (GB)	NVMe read caching (GB)	Disk throughpu t (MBps)		SSD drive IOPS *		
	Baseline	Burst			Baseline	Burst	Baseline	Burst	
128	188	1,500	Tens	16	–	128	1,250	6,000	40,000
256	375	1,500	of thousands baseline	32	–	256	1,250	12,000	40,000
512	750	1,500	Hundreds	64	–	512	1,250	20,000	40,000
1,024	1,500	–	of thousands baseline	128	–	1,024	1,250	40,000	–
2,048	3,125	–		256	1,900	2,048	–	80,000	–
4,096	6,250	–		512	5,400	4,096	–	160,000	–

Note

* Your SSD IOPS are only used when you access data that is not cached in your file server's in-memory cache or NVMe cache.

These performance specifications apply to Single-AZ scale-up file systems in all other Amazon Web Services Regions where FSx for ONTAP is available.

Performance specifications for file systems in [all other Amazon Web Services Regions where FSx for ONTAP is available](#)

FSx throughput capacity (MBps)	Network throughput capacity (MBps)		Network IOPS	In-memory caching (GB)	Disk throughput (MBps)		SSD drive IOPS *	
	Baseline	Burst			Baseline	Burst	Baseline	Burst
128	150	1,250	Tens of thousands baseline	16	128	600	6,000	18,750
256	300	1,250	baseline	32	256	600	12,000	18,750
512	625	1,250	Hundreds of thousands baseline	64	512	600	18,750	–
1,024	1,500	–	of thousands baseline	128	1,024	–	40,000	–
2,048	3,125	–	baseline	256	2,048	–	80,000	–

Note

* Your SSD IOPS are only used when you access data that is not cached in your file server's in-memory cache or NVMe cache.

Single-AZ (scale-out)

These performance specifications apply to scale-out file systems.

Performance specifications for scale-out file systems

FSx throughput capacity (MBps)	Network throughput capacity (MBps)		Network IOPS	In- memory caching (GB)	Disk throughput (MBps)		SSD drive IOPS *	
	Baseline	Burst			Baseline	Burst	Baseline	Burst
3,072**	6,250	–	Hundreds of thousands baseline	128	3,072	–	100,000	–
6,144**	12,500	–		256	6,144	–	200,000	–

Note

* Your SSD IOPS are only used when you access data that is not cached in your file server's in-memory cache or NVMe cache.

** Per HA pair (up to 12). For more information, see [High-availability \(HA\) pairs](#).

Multi-AZ (scale-up)

These performance specifications apply to Multi-AZ scale-up file systems created after November 28, 2022 in the specified Amazon Web Services Regions.

Performance specifications for file systems in the following Amazon Web Services Regions: US East (N. Virginia), US East (Ohio), US West (Oregon), and Europe (Ireland)

FSx throughput capacity (MBps)	Network throughput capacity (MBps)		Network IOPS	In- memory caching (GB)	NVMe caching (GB)	Disk throughput (MBps)		SSD drive IOPS *	
	Baseline	Burst				Baseline	Burst	Baseline	Burst

FSx throughput capacity (MBps)	Network throughput capacity (MBps)	Network IOPS	In-memory caching (GB)	NVMe caching (GB)	Disk throughput (MBps)	SSD drive IOPS *			
128	188	1,500	Tens of thousands baseline	16	238	128	1,250	6,000	40,000
256	375	1,500	32	475	256	1,250	12,000	40,000	
512	750	1,500	Hundreds of thousands baseline	64	950	512	1,250	20,000	40,000
1,024	1,500	–	128	1,900	1,024	1,250	40,000	–	
2,048	3,125	–	256	3,800	2,048	–	80,000	–	
4,096	6,250	–	512	7,600	4,096	–	160,000	–	

Note

* Your SSD IOPS are only used when you access data that is not cached in your file server's in-memory cache or NVMe cache.

These performance specifications apply to Multi-AZ scale-up file systems in all other Amazon Web Services Regions where FSx for ONTAP is available.

Performance specifications for file systems in [all other Amazon Web Services Regions where FSx for ONTAP is available](#)

FSx throughput capacity (MBps)	Network throughput capacity (MBps)	Network IOPS	In-memory caching (GB)	NVMe caching (GB)	Disk throughput (MBps)	SSD drive IOPS *
	Baseline	Burst			Baseline	Burst
						Baseline
						Burst

FSx throughput capacity (MBps)	Network throughput capacity (MBps)	Network IOPS	Network In-memory caching (GB)	NVMe caching (GB)	Disk throughput (MBps)	SSD drive IOPS *			
128	150	1,250	Tens of thousands baseline	16	150	128	600	6,000	18,750
256	300	1,250	32	300	256	600	600	12,000	18,750
512	625	1,250	Hundreds of thousands baseline	64	600	512	600	18,750	–
1,024	1,500	–	128	1,200	1,024	–	–	40,000	–
2,048	3,125	–	256	2,400	2,048	–	–	80,000	–

Note

* Your SSD IOPS are only used when you access data that is not cached in your file server's in-memory cache or NVMe cache.

Example: storage capacity and throughput capacity

The following example illustrates how storage capacity and throughput capacity impact file system performance.

A scale-up file system that is configured with 2 TiB of SSD storage capacity and 512 MBps of throughput capacity has the following throughput levels:

- Network throughput – 625 MBps baseline and 1,250 MBps burst (see throughput capacity table)
- Disk throughput – 512 MBps baseline and 600 MBps burst.

Your workload accessing the file system will therefore be able to drive up to 625 MBps baseline and 1,250 MBps burst throughput for file operations performed on actively accessed data cached in the file server in-memory cache and NVMe cache.

Administering FSx for ONTAP resources

Using the Amazon Web Services Management Console, Amazon CLI, and ONTAP CLI and API, you can perform the following administrative actions for FSx for ONTAP resources:

- Creating, listing, updating, and deleting file systems, storage virtual machines (SVMs), volumes, backups, and tags.
- Managing access, administrative accounts and passwords, password requirements, SMB and iSCSI protocols, network accessibility for the mount targets of existing file systems

Topics

- [Managing FSx for ONTAP file systems](#)
- [Creating FSx for ONTAP file systems](#)
- [Updating a file system](#)
- [Deleting a file system](#)
- [Viewing file system details](#)
- [Managing FSx for ONTAP storage virtual machines](#)
- [Managing FSx for ONTAP volumes](#)
- [Creating an iSCSI LUN](#)
- [Managing SMB shares](#)
- [File access auditing](#)
- [Scaling SSD storage capacity and provisioned IOPS](#)
- [Managing throughput capacity](#)
- [Optimizing performance with Amazon FSx maintenance windows](#)
- [Tag your Amazon FSx resources](#)
- [Managing FSx for ONTAP resources using NetApp applications](#)

Managing FSx for ONTAP file systems

A file system is the primary Amazon FSx resource, analogous to an on-premises ONTAP cluster. You specify the solid state drive (SSD) storage capacity and throughput capacity for your file system, and choose a virtual private cloud (VPC) in which to create the file system. Each file system has

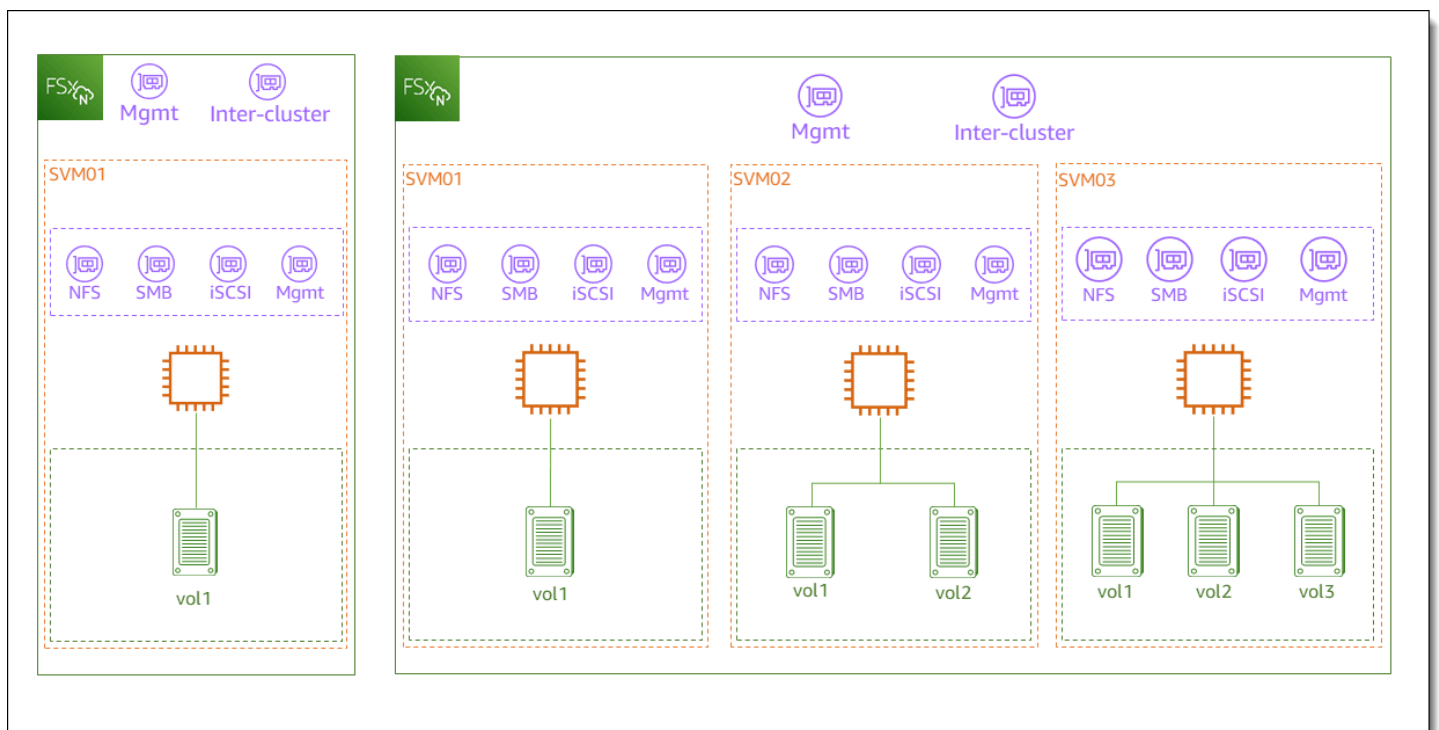
a management endpoint that you can use to manage resources and data with the ONTAP CLI or REST API.

File system resources

An Amazon FSx for NetApp ONTAP file system is composed of the following primary resources:

- The physical hardware of the file system itself, which includes the file servers and storage media.
- One or more highly-available (HA) file server pairs, which host your storage virtual machines (SVMs). Scale-up file systems have one HA pair, and scale-out file systems have two or more HA pairs. Each HA pair has a storage pool called an aggregate. The collection of aggregates across all HA pairs makes up your SSD storage tier.
- One or more storage virtual machines (SVMs) that host the file system volumes and have their own credentials and access management.
- One or more volumes that virtually organize your data and are mounted by your clients.

The following image illustrates the architecture of a scale-up FSx for ONTAP file system with one HA pair, and the relationship between its primary resources. The FSx for ONTAP file system on the left is the simplest file system, with one SVM and one volume. The file system on the right has multiple SVMs, with some SVMs having multiple volumes. File systems and SVMs each have multiple management endpoints, and SVMs also have data access endpoints.



When creating an FSx for ONTAP file system, you define the following properties:

- **Deployment type** – The deployment type of your file system (Multi-AZ or Single-AZ). Single-AZ file systems replicate your data and offer automatic failover within a single Availability Zone, and offer scale-out file systems. Multi-AZ file systems provide added resiliency by also replicating your data and supporting failover across multiple Availability Zones within the same Amazon Web Services Region.
- **Storage capacity** – This is the amount of SSD storage, up to 192 tebibytes (TiB) for scale-up file systems and 1 pebibyte (PiB) for scale-out file systems.
- **SSD IOPS** – By default, each gigabyte of SSD storage includes three SSD IOPS (up to the maximum supported by your file system configuration). You can optionally provision additional SSD IOPS as needed.
- **Throughput capacity** – The sustained speed at which the file server can serve data.
- **Networking** – The VPC and subnets for the management and data access endpoints that your file system creates. For a Multi-AZ file system, you also define an IP address range and route tables.
- **Encryption** – The Amazon Key Management Service (Amazon KMS) key that's used to encrypt the file system data at rest.
- **Administrative access** – You can specify the password for the `fsxadmin` user. You can use this user to administer the file system by using the NetApp ONTAP CLI and REST API.

You can manage FSx for ONTAP file systems by using the NetApp ONTAP CLI or REST API. You can also set up SnapMirror or SnapVault relationships between an Amazon FSx file system and another ONTAP deployment (including another Amazon FSx file system). Each FSx for ONTAP file system has the following file system endpoints that provide access to NetApp applications:

- **Management** – Use this endpoint to access the NetApp ONTAP CLI over Secure Shell (SSH), or to use the NetApp ONTAP REST API with your file system.
- **Intercluster** – Use this endpoint when setting up replication using NetApp SnapMirror or caching using NetApp FlexCache.

For more information, see [Managing FSx for ONTAP resources using NetApp applications](#) and [Scheduled replication using NetApp SnapMirror](#).

High-availability (HA) pairs

Each FSx for ONTAP file system is powered by one or multiple high-availability (HA) pairs of file servers in an active-standby configuration. In this configuration, there is a preferred file server that actively serves traffic and a secondary file server that takes over if the active server is unavailable. FSx for ONTAP scale-up file systems are powered by one HA pair, which delivers up to 4 GBps of throughput capacity and 160,000 SSD IOPs. FSx for ONTAP scale-out file systems are powered by up to 12 HA pairs, which can deliver up to 72 GBps of throughput capacity and 2,400,000 SSD IOPS (6 GBps of throughput capacity and 200,000 SSD IOPS per HA pair).

When you create your file system from the Amazon FSx console, Amazon FSx recommends the number of HA pairs that you should use based on your desired SSD storage. You can also manually choose the number of HA pairs based on your workload and performance requirements. We recommend that you use a single HA pair if your file system requirements are satisfied by up to 4 GBps of throughput capacity and 160,000 SSD IOPs, and multiple HA pairs if your workloads need higher levels of performance scalability.

Each HA pair has one aggregate, which is a logical set of physical disks.

Note

You can't add HA pairs to existing file systems. Instead, you can migrate data between file systems (with different HA pairs) using SnapMirror, Amazon DataSync, or by restoring your data from a backup to a new file system.

Creating FSx for ONTAP file systems

This section describes how to create an FSx for ONTAP file system using the Amazon FSx console, Amazon CLI, or the Amazon FSx API. You can create a file system in a virtual private cloud (VPC) that you own, or in a VPC that another Amazon Web Services account has shared with you. There are considerations when creating a Multi-AZ file system in a VPC in which you are a participant. These considerations are explained in this topic.

By default, when you create a new file system from the Amazon FSx console, Amazon FSx automatically creates a file system with a single storage virtual machine (SVM) and one volume, allowing for quick access to data from Linux instances over the Network File System (NFS) protocol. When creating the file system, you can optionally join the SVM to an Active Directory to enable

access from Windows and macOS clients over the Server Message Block (SMB) protocol. After your file system is created, you can create additional SVMs and volumes as needed.

To create a file system (console)

This procedure uses the **Standard create** creation option to create an FSx for ONTAP file system with a configuration that you customize for your needs. For information about using the **Quick create** creation option to rapidly create a file system with a default set of configuration parameters, see [Step 1: Create an Amazon FSx for NetApp ONTAP file system](#).

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. On the dashboard, choose **Create file system**.
3. On the **Select file system type** page, for **File system options**, choose **Amazon FSx for NetApp ONTAP**, and then choose **Next**.
4. In the **Creation method** section, choose **Standard create**.
5. In the **File system details** section, provide the following information:
 - For **File system name - optional**, enter a name for your file system. It's easier to find and manage your file systems when you name them. You can use a maximum of 256 Unicode letters, white space, and numbers, plus these special characters: + - = . _ : /
 - For **Deployment type** choose **Multi-AZ** or **Single-AZ**.
 - **Multi-AZ** file systems replicate your data and support failover across multiple Availability Zones in the same Amazon Web Services Region.
 - **Single-AZ** file systems replicate your data and offer automatic failover within a single Availability Zone.

Note

Choose **Single-AZ** if you want the option to create a file system with two or more high-availability (HA) pairs (up to 12). For more information, see [High-availability \(HA\) pairs](#).

For more information, see [Availability and durability](#).

- For **SSD storage capacity**, enter the storage capacity of your file system, in gibibytes (GiB). Enter any whole number in the range of 1,024–1,048,576 GiB (up to 1 pebibyte [PiB]).

You can increase the amount of storage capacity as needed at any time after you create the file system. For more information, see [Managing storage capacity](#).

- For **Provisioned SSD IOPS**, you have two options to provision the number of IOPS for your file system:
 - Choose **Automatic** (the default) if you want Amazon FSx to automatically provision 3 IOPS per GiB of SSD storage.
 - Choose **User-provisioned** if you want to specify the number of IOPS. You can provision a maximum of 200,000 SSD IOPS per file system.

Note

You can increase your provisioned SSD IOPS after you create the file system. Keep in mind that the maximum level of SSD IOPS your file system can achieve is also dictated by your file system's throughput capacity even when provisioning additional SSD IOPS. For more information, see [Impact of throughput capacity on performance](#) and [Managing storage capacity](#).


- For **Throughput capacity**, you have two options for determining your throughput capacity in megabytes per second (MBps):
 - Choose **Recommended throughput capacity** if you want Amazon FSx to automatically choose the throughput capacity based on the amount of storage capacity that you chose.
 - Choose **Specify throughput capacity** if you want to specify the amount of throughput capacity. If you choose this option, a **Throughput capacity** dropdown appears and is populated based on the deployment type that you chose. You can also choose the number of HA pairs (up to 12). For more information, see [High-availability \(HA\) pairs](#).

Throughput capacity is the sustained speed at which the file server that hosts your file system can serve data. For more information, see [Amazon FSx for NetApp ONTAP performance](#).

6. In the **Networking** section, provide the following information:

- For **Virtual Private Cloud (VPC)**, choose the VPC that you want to associate with your file system.
- For **VPC Security Groups**, you can choose a security group to associate with your file system's network interface. If you don't specify one, Amazon FSx will associate the VPC's default security group with your file system.

- Specify a **Subnet** for your file server. If you are creating a Multi-AZ file system, also choose a **Standby subnet** for the standby file server.
- (Multi-AZ only) For **VPC route tables**, specify the VPC route tables to create your file system's endpoints. Select all VPC route tables associated with the subnets in which your clients are located. By default, Amazon FSx selects your VPC's default route table. For more information, see [Accessing data from outside the deployment VPC](#).


 **Note**

Amazon FSx manages these route tables for Multi-AZ file systems using tag-based authentication. These route tables are tagged with Key: AmazonFSx; Value: ManagedByAmazonFSx. When creating FSx for ONTAP Multi-AZ file systems using Amazon CloudFormation we recommend that you add the Key: AmazonFSx; Value: ManagedByAmazonFSx tag manually.

- (Multi-AZ only) **Endpoint IP address range** specifies the IP address range in which the endpoints to access your file system are created.

You have three options for the endpoint IP address range:

- **Unallocated IP address range from your VPC** – Amazon FSx chooses the last 64 IP addresses from the VPC's primary CIDR range to use as the endpoint IP address range for the file system. This range is shared across multiple file systems if you choose this option multiple times.

 **Note**


This option is grayed out if any of the last 64 IP addresses in a VPC's primary CIDR range are in use by a subnet. In this case, you can still choose an in-VPC address range (that is, a range that's not at the end of your primary CIDR range or a range that's in a secondary CIDR of your VPC) by choosing the **Enter an IP address range** option.

- For **Preferred subnet**, specify a **Subnet** for your file server. If you are creating a Multi-AZ file system, also choose a **Standby subnet** for the standby file server.
- (Multi-AZ only) For **VPC route tables**, specify the VPC route tables to create your file system's endpoints. Select all VPC route tables associated with the subnets in which your clients are located. By default, Amazon FSx selects your VPC's default route table.

- (Multi-AZ only) **Endpoint IP address range** specifies the IP address range in which the endpoints to access your file system are created.

You have three options for the endpoint IP address range:

- **Unallocated IP address range from your VPC** – Amazon FSx chooses the last 64 IP addresses from the VPC's primary CIDR range to use as the endpoint IP address range for the file system. This range is shared across multiple file systems if you choose this option multiple times.

 **Note**

This option is grayed out if any of the last 64 IP addresses in a VPC's primary CIDR range are in use by a subnet. In this case, you can still choose an in-VPC address range (that is, a range that's not at the end of your primary CIDR range or a range that's in a secondary CIDR of your VPC) by choosing the **Enter an IP address range** option.

- **Floating IP address range outside your VPC** – Amazon FSx chooses a 198.19.x.0/24 address range that isn't already used by any other file systems with the same VPC and route tables.
- **Enter an IP address range** – You can provide a CIDR range of your own choosing. The IP address range that you choose can either be inside or outside the VPC's IP address range, as long as it doesn't overlap with any subnet.

 **Note**

Do not choose any range that falls within the following CIDR ranges, as they are incompatible with FSx for ONTAP:

- 0.0.0.0/8
- 127.0.0.0/8
- 198.19.0.0/20
- 224.0.0.0/4
- 240.0.0.0/4
- 255.255.255.255/32

7. In the **Security & encryption** section, for **Encryption key**, choose the Amazon Key Management Service (Amazon KMS) encryption key that protects your file system's data at rest.
8. For **File system administrative password**, enter a secure password for the fsxadmin user. Confirm the password.

You can use the fsxadmin user to administer your file system using the ONTAP CLI and REST API. For more information about the fsxadmin user, see [Managing file systems with the ONTAP CLI](#).

9. In the **Default storage virtual machine configuration** section, provide the following information:
 - In the **Storage virtual machine name** field, provide a name for the storage virtual machine. You can use a maximum of 47 alphanumeric characters, plus the underscore (_) special character.
 - For **SVM administrative password**, you can optionally choose **Specify a password** and provide a password for the SVM's vsadmin user. You can use the vsadmin user to administer the SVM using the ONTAP CLI or REST API. For more information about the vsadmin user, see [Managing SVMs with the ONTAP CLI](#).

If you choose **Don't specify a password** (the default), you can still use the file system's fsxadmin user to manage your file system using the ONTAP CLI or REST API, but you can't use your SVM's vsadmin user to do the same.

- In the **Active Directory** section, you can join an Active Directory to the SVM. For more information, see [Working with Microsoft Active Directory in FSx for ONTAP](#).

If you don't want to join your SVM to an Active Directory, choose **Do not join an Active Directory**.

If you want to join your SVM to a self-managed Active Directory domain, choose **Join an Active Directory**, and provide the following details for your Active Directory:

- The NetBIOS name of the Active Directory computer object to create for your SVM. The NetBIOS name cannot exceed 15 characters.
- The fully qualified domain name of your Active Directory. The domain name cannot exceed 255 characters.
- **DNS server IP addresses** – The IPv4 addresses of the Domain Name System (DNS) servers for your domain.

- **Service account username** – The user name of the service account in your existing Active Directory. Do not include a domain prefix or suffix.
- **Service account password** – The password for the service account.
- **Confirm password** – The password for the service account.
- (Optional) **Organizational Unit (OU)** – The distinguished path name of the organizational unit to which you want to join your file system.
- **Delegated file system administrators group** – The name of the group in your Active Directory that can administer your file system.

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

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

10. In the **Default volume configuration** section, provide the following information for the default volume that is created with your file system:

- In the **Volume name** field, provide a name for the volume. You can use up to 203 alphanumeric or underscore (_) characters.
- (*Scale-up file systems only*) For **Volume style**, choose either **FlexVol** or **FlexGroup**. FlexVol volumes are general-purpose volumes that can be up to 300 TiB in size. FlexGroup volumes are intended for high-performance workloads and can be up to 20 PiB in size.
- For **Volume size**, enter any whole number in the range of 800 gibibytes (GiB)–2,000 pebibytes (PiB).
- For **Volume type**, choose **Read-Write (RW)** to create a volume that is readable and writable or **Data Protection (DP)** to create a volume that is read-only and can be used as the destination of a NetApp SnapMirror or SnapVault relationship. For more information, see [Volume types](#).
- For **Junction path**, enter a location within the file system to mount the volume. The name must have a leading forward slash, for example /vol3.
- For **Storage efficiency**, choose **Enabled** to enable the ONTAP storage-efficiency features (deduplication, compression, and compaction). For more information, see [FSx for ONTAP storage efficiency](#).

- For **Volume security style**, choose between **Unix (Linux)**, **NTFS**, and **Mixed** for the volume. For more information, see [Volume security style](#).
- For **Snapshot policy**, choose a snapshot policy for the volume. For more information about snapshot policies, see [Snapshot policies](#).

If you choose **Custom policy**, you must specify the policy's name in the **custom-policy** field. The custom policy must already exist on the SVM or in the file system. You can create a custom snapshot policy with the ONTAP CLI or REST API. For more information, see [Create a Snapshot Policy](#) in the NetApp ONTAP Product Documentation.

11. In the **Default volume storage tiering** section, for **Capacity pool tiering policy**, choose the storage pool tiering policy for the volume, which can be **Auto** (the default), **Snapshot Only**, **All**, or **None**. For more information about capacity pool tiering policies, see [Volume tiering policies](#).

For **Tiering policy cooling period**, if you have set storage tiering to either Auto and Snapshot-only policies, valid values are 2-183 days. A volume's tiering policy cooling period defines the number of days before data that has not been accessed is marked cold and moved to capacity pool storage.

12. In **Backup and maintenance - optional**, you can set the following options:

- For **Daily automatic backup**, choose **Enabled** for automatic daily backups. This option is enabled by default.
- For **Daily automatic backup window**, set the time of the day in Coordinated Universal Time (UTC) that you want the daily automatic backup window to start. The window is 30 minutes starting from this specified time. This window can't overlap with the weekly maintenance backup window.
- For **Automatic backup retention period**, set a period from 1–90 days that you want to retain automatic backups.
- For **Weekly maintenance window**, you can set the time of the week that you want the maintenance window to start. Day 1 is Monday, 2 is Tuesday, and so on. The window is 30 minutes starting from this specified time. This window can't overlap with the daily automatic backup window.

13. For **Tags - optional**, you can enter a key and value to add tags to your file system. A tag is a case-sensitive key-value pair that helps you manage, filter, and search for your file system.

Choose **Next**.

14. Review the file system configuration shown on the **Create file system** page. For your reference, note which file system settings you can modify after the file system is created.
15. Choose **Create file system**.

To create a file system (CLI)

- To create an FSx for ONTAP file system, use the [create-file-system](#) CLI command (or the equivalent [CreateFileSystem](#) API operation), as shown in the following example.

```
aws fsx create-file-system \
  --file-system-type ONTAP \
  --storage-capacity 1024 \
  --storage-type SSD \
  --security-group-ids security-group-id \

  --subnet-ids subnet-abcdef1234567890b subnet-abcdef1234567890c \
  --ontap-configuration DeploymentType=MULTI_AZ_1,
  ThroughputCapacity=512,PreferredSubnetId=subnet-abcdef1234567890b
```

After successfully creating the file system, Amazon FSx returns the file system's description in JSON format as shown in the following example.

```
{
  "FileSystem": {
    "OwnerId": "111122223333",
    "CreationTime": 1625066825.306,
    "FileSystemId": "fs-0123456789abcdef0",
    "FileSystemType": "ONTAP",
    "Lifecycle": "CREATING",
    "StorageCapacity": 1024,
    "StorageType": "SSD",
    "VpcId": "vpc-11223344556677aab",
    "SubnetIds": [
      "subnet-abcdef1234567890b",
      "subnet-abcdef1234567890c"
    ],
    "KmsKeyId": "arn:aws:kms:us-east-1:111122223333:key/wJa1rXUtnFEMI/K7MDENG/
bPxRfiCYEXAMPLEKEY",
    "ResourceARN": "arn:aws:fsx:us-east-1:111122223333:file-system/
fs-0123456789abcdef0",
```

```

"Tags": [],
"OntapConfiguration": {
  "DeploymentType": "MULTI_AZ_HA_1",
  "EndpointIpAddressRange": "198.19.0.0/24",
  "Endpoints": {
    "Management": {
      "DnsName": "management.fs-0123456789abcdef0.fsx.us-east-1.amazonaws.com"
    },
    "Intercluster": {
      "DnsName": "intercluster.fs-0123456789abcdef0.fsx.us-east-1.amazonaws.com"
    }
  },
  "DiskIopsConfiguration": {
    "Mode": "AUTOMATIC",
    "Iops": 3072
  },
  "PreferredSubnetId": "subnet-abcdef1234567890b",
  "RouteTableIds": [
    "rtb-abcdef1234567890e",
    "rtb-abcd1234ef567890b"
  ],
  "ThroughputCapacity": 512,
  "WeeklyMaintenanceStartTime": "4:10:00"
}
}
}

```

Note

Unlike the process of creating a file system in the console, the `create-file-system` CLI command and the `CreateFileSystem` API operation don't create a default SVM or volume. To create an SVM, see [Creating a storage virtual machine](#); to create a volume, see [Creating volumes](#).

Creating FSx for ONTAP file systems in shared subnets

VPC sharing enables multiple Amazon Web Services accounts to create resources into shared, centrally-managed virtual private clouds (VPCs). In this model, the account that owns the VPC (owner) shares one or more subnets with other accounts (participants) that belong to the same organization from Amazon Organizations.

Participant accounts can create FSx for ONTAP Single-AZ and Multi-AZ file systems in a VPC subnet that the owner account has shared with them. For a participant account to create a Multi-AZ file system, the owner account also needs to grant Amazon FSx permission to modify route tables in the shared subnets on behalf of the participant account. For more information, see [Managing shared VPC support for Multi-AZ file systems](#).

Note

It is the participant account's responsibility to coordinate with the VPC owner to prevent the creation of any subsequent VPC subnets that will overlap with the in-VPC CIDR of the participant's file systems. If subnets do overlap, traffic to the file system can get interrupted.

Shared subnet requirements and considerations

When creating FSx for ONTAP file systems into shared subnets, note the following:

- The owner of the VPC subnet must share a subnet with a participant account before that account can create an FSx for ONTAP file system in it.
- You can't launch resources using the default security group for the VPC because it belongs to the owner. Additionally, participant accounts can't launch resources using security groups that are owned by other participants or the owner.
- In a shared subnet, the participant and the owner separately controls the security groups within each respective account. The owner account can see security groups that are created by the participants, but cannot perform any actions on them. If the owner account wants to remove or modify these security groups, the participant that created the security group must take the action.
- Participant accounts can view, create, modify, and delete Single-AZ file systems and their associated resources in subnets that the owner account has shared with them.
- Participant accounts can create, view, modify, and delete Multi-AZ file systems and their associated resources in subnets that the owner account has shared with them. Additionally, the owner account must also grant the Amazon FSx service permissions to modify route tables in the shared subnets on behalf of the participants account. For more information, see [Managing shared VPC support for Multi-AZ file systems](#)
- The shared VPC owner cannot view, modify, or delete resources that a participant creates in the shared subnet. This is in addition to the VPC resources that each account has different access to.

For more information, see [Responsibilities and permissions for owners and participants](#) in the Amazon VPC User Guide.

For more information, see [Share your VPC with other accounts](#) in the Amazon VPC User Guide.

When sharing a VPC subnet

When sharing your subnets with participant accounts that will be creating FSx for ONTAP file systems in the shared subnets, you will need to do the following:

- The VPC owner needs to use Amazon Resource Access Manager to securely share VPCs and subnets with other Amazon Web Services accounts. For more information, see [Sharing your Amazon resources](#) in the Amazon Resource Access Manager User Guide.
- The VPC owner needs to share one or more VPCs with a participant account. For more information, see [Share your VPC with other accounts](#) in the Amazon Virtual Private Cloud User Guide.
- For participant accounts to create FSx for ONTAP Multi-AZ file systems, the VPC owner must also grant the Amazon FSx service permissions to create and modify route tables in the shared subnets on behalf of the participant accounts. This is because FSx for ONTAP Multi-AZ file systems use floating IP addresses so that connected clients can seamlessly transition between the preferred and standby file servers during a failover event. When a failover event occurs, Amazon FSx updates all routes in all route tables associated with the file system to point to the currently active file server.

Managing shared VPC support for Multi-AZ file systems

Owner accounts can manage whether or not participant accounts can create Multi-AZ FSx for ONTAP file systems in VPC subnets that the owner has shared with participants using the Amazon Web Services Management Console, Amazon CLI, and API, as described in the following sections.

To manage VPC sharing for Multi-AZ file systems (console)

Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.

1. In the navigation pane, choose **Settings**.
2. Locate the **Multi-AZ shared VPC settings** on the **Settings** page.

- To enable VPC sharing for Multi-AZ file systems in VPC subnets that you share, choose **Enable route table updates from participant accounts**.
- To disable VPC sharing for Multi-AZ file systems in all VPCs that you own, choose **Disable route table updates from participant accounts**. The confirmation screen is displayed.

⚠ Important

We strongly recommend that participant-created Multi-AZ file systems in the shared VPC are deleted before you disable this feature. Once the feature is disabled, these file systems will enter a MISCONFIGURED state and will be at risk of becoming unavailable.

3. Enter **confirm** and choose **Confirm** to disable the feature.

To manage VPC sharing for Multi-AZ file systems (Amazon CLI)

1. To view the current setting for Multi-AZ VPC sharing, use the [describe-shared-vpc-configuration](#) CLI command, or the equivalent [DescribeSharedVpcConfiguration](#) API command, shown as follows:

```
$ aws fsx describe-shared-vpc-configuration
```

The service responds to a successful request as follows:

```
{
  "EnableFsxRouteTableUpdatesFromParticipantAccounts": "false"
}
```

2. To manage the Multi-AZ shared VPC configuration, use the [update-shared-vpc-configuration](#) CLI command, or the equivalent [UpdateSharedVpcConfiguration](#) API command. The following example enables VPC sharing for Multi-AZ file systems.

```
$ aws fsx update-shared-vpc-configuration --enable-fsx-route-table-updates-from-participant-accounts true
```

The service responds to a successful request as follows:

```
{
```

```
"EnableFsxRouteTableUpdatesFromParticipantAccounts": "true"
}
```

- To disable the feature, set `EnableFsxRouteTableUpdatesFromParticipantAccounts` to `false`, as shown in the following example.

```
$ aws fsx update-shared-vpc-configuration --enable-fsx-route-table-updates-from-participant-accounts false
```

The service responds to a successful request as follows:

```
{
  "EnableFsxRouteTableUpdatesFromParticipantAccounts": "false"
}
```

Updating a file system

This topic explains which properties of an existing file system that you can update, and provides procedures to do so using the console and CLI.

You can update the following FSx for ONTAP file system properties using the Amazon FSx console, the Amazon CLI, and the Amazon FSx API:

- Automatic daily backups.** Turns automatic daily backups on or off, modifies the backup window and the backup retention period. For more information about backups, see [Working with automatic daily backups](#).
- Weekly maintenance window.** Sets the day of the week and time that Amazon FSx performs file system maintenance and updates. For more information about maintenance window, see [Optimizing performance with Amazon FSx maintenance windows](#).
- File system administrative password.** Changes the password for the file system's `fsxadmin` user. You can use the `fsxadmin` user to administer your file system using the ONTAP CLI and REST API. For more information about the `fsxadmin` user, see [Managing file systems with the ONTAP CLI](#).
- Amazon VPC route tables.** With Multi-AZ FSx for ONTAP file systems, the endpoints you use to access data over NFS or SMB and the management endpoints to access the ONTAP CLI, API, and BlueXP use floating IP addresses in the Amazon VPC route tables that you associate with your file system. You can associate new route tables that you create with your existing Multi-AZ file

systems—allowing you to configure which clients can access your data even as your network evolves. You can also disassociate (remove) existing route tables from your file system.

Note

Amazon FSx manages VPC route tables for Multi-AZ file systems using tag-based authentication. These route tables are tagged with Key: AmazonFSx; Value: ManagedByAmazonFSx. When creating or updating FSx for ONTAP Multi-AZ file systems using Amazon CloudFormation we recommend that you add the Key: AmazonFSx; Value: ManagedByAmazonFSx tag manually.

To update a file system (console)

The following procedures provide you with instructions on how to make updates to an existing FSx for ONTAP file system using the Amazon Web Services Management Console.

To update automatic daily backups

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. To display the file system details page, in the left navigation pane, choose **File systems**, and then choose the FSx for ONTAP file system that you want to update.
3. Choose the **Backups** tab in the second panel on the page.
4. Choose **Update**.
5. Modify the automatic daily backup settings for this file system.
6. Choose **Save** to save your changes.

To update the weekly maintenance window

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. To display the file system details page, in the left navigation pane, choose **File systems**, and then choose the FSx for ONTAP file system that you want to update.
3. Choose the **Administration** tab in the second panel on the page.
4. In the **Maintenance** pane, choose **Update**.
5. Modify when the weekly maintenance window occurs for this file system.
6. Choose **Save** to save your changes.

To change the file system administrative password

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. To display the file system details page, in the left navigation pane, choose **File systems**, and then choose the FSx for ONTAP file system that you want to update.
3. Choose the **Administration** tab.
4. In the **ONTAP administration** pane, choose **Update** under **ONTAP administrator password**.
5. In the **Update ONTAP administrator credentials** dialog box, enter a new password in the **ONTAP administrative password** field.
6. Use the **Confirm password** field to confirm the password.
7. Choose **Update credentials** to save your change.

Note

If you receive an error stating that the new password does not meet the password requirements, you can use the [security login role config show](#) ONTAP CLI command to view the password requirement settings on the file system. For more information, including instructions on how to change password setting, see [Updating the fsxadmin account password fails](#).

To update VPC route tables on Multi-AZ file systems

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. To display the file system details page, in the left navigation pane, choose **File systems**, and then choose the FSx for ONTAP file system that you want to update.
3. For **Actions**, choose **Manage route tables**. This option is only available for Multi-AZ file systems.
4. In the **Manage route tables** dialog box, do one of the following:
 - To associate a new VPC route table, select a route table from the **Associate new route tables** dropdown list, and then choose **Associate**.
 - To disassociate an existing VPC route table, select a route table from the **Current route tables** pane, and then choose **Disassociate**.
5. Choose **Close**.

To update a file system (CLI)

The following procedure illustrates how to make updates to an existing FSx for ONTAP file system using the Amazon CLI.

1. To update the configuration of an FSx for ONTAP file system, use the [update-file-system](#) CLI command (or the equivalent [UpdateFileSystem](#) API operation), as shown in the following example.

```
aws fsx update-file-system \  
  --file-system-id fs-0123456789abcdef0 \  
  --ontap-configuration  
    AutomaticBackupRetentionDays=30,DailyAutomaticBackupStartTime=01:00, \  
    WeeklyMaintenanceStartTime=1:01:30,AddRouteTableIds=rtb-0123abcd, \  
    FsxAdminPassword=new-fsx-admin-password
```

2. To disable automatic daily backups, set the `AutomaticBackupRetentionDays` property to 0.

```
aws fsx update-file-system \  
  --file-system-id fs-0123456789abcdef0 \  
  --ontap-configuration AutomaticBackupRetentionDays=0
```

Deleting a file system

You can delete an FSx for ONTAP file system using the Amazon FSx console, the Amazon CLI, and the Amazon FSx API and SDKs.

To delete a file system:

- **Using the console** – Follow the procedure described in [Step 3: Clean up resources](#).
- **Using the CLI or API** – First delete all the volumes and SVMs on your file system. Then use the [delete-file-system](#) CLI command or the [DeleteFileSystem](#) API operation.

Viewing file system details

You can view detailed configuration information for your FSx for ONTAP file system using the Amazon FSx console, the Amazon CLI, and the API and supported Amazon SDKs.

To view detailed file system information:

- **Using the console** – Choose a file system to view the **File systems** detail page. The **Summary** panel shows the file system's ID, life cycle status, deployment type, SSD storage capacity, throughput capacity, provisioned IOPS, Availability Zones, and creation time.

The following tabs provide detailed configuration information and editing for properties that can be modified:

- Network & security
- Monitoring & performance – Displays CloudWatch alarms you've created, and metrics and warnings for the following categories:
 - Summary – high level summary of file system activity metrics
 - File system storage capacity
 - File server and disk performance

For more information, see [Monitoring with Amazon CloudWatch](#).

- Administration – Displays the following file system administration information:
 - The DNS names and IP addresses of the file system's management and inter-cluster endpoints.
 - The ONTAP administrator username.
 - The option to update the ONTAP administrator password.
 - List of the file system's SVMs
 - List of the file system's volumes
 - Backup settings – change the file system's automatic daily backup setting.
 - Updates – shows the status of user initiated updates made to the file system's configuration.
 - Tags – view, edit, add, remove tag Key:Value pairs.
- **Using the CLI or API** – Use the [describe-file-systems](#) CLI command or the [DescribeFileSystems](#) API operation.

FSx for ONTAP file system status

You can view the status of an Amazon FSx file system by using the Amazon FSx console, the Amazon CLI command [describe-file-systems](#), or the API operation [DescribeFileSystems](#).

File system status	Description
AVAILABLE	The file system has been successfully created and is available for use.
CREATING	Amazon FSx is creating a new file system.
DELETING	Amazon FSx is deleting an existing file system.
MISCONFIGURED	The file system is in a misconfigured but recoverable state.
FAILED	<ol style="list-style-type: none"> 1. The file system has failed and Amazon FSx can't recover it. 2. When creating new file system, Amazon FSx was unable to create a new file system.

Managing FSx for ONTAP storage virtual machines

In FSx for ONTAP, volumes are hosted on virtual file servers called storage virtual machines (SVMs). An SVM is an isolated file server with its own administrative credentials and endpoints for administering and accessing data. When you access data in FSx for ONTAP, your clients and workstations mount a volume, SMB share, or iSCSI LUN hosted by an SVM using the SVM's endpoint (IP address).

Amazon FSx automatically creates a default SVM on your file system when you create a file system using the Amazon Web Services Management Console. You can create additional SVMs on your file system at any time using the console, Amazon CLI, or Amazon FSx API and SDKs. You cannot create SVMs using the ONTAP CLI or REST API.

You can join your SVMs to a Microsoft Active Directory for file access authentication and authorization. For more information, see [Working with Microsoft Active Directory in FSx for ONTAP](#).

Maximum number of SVMs per file system

The following table lists the maximum number of SVMs that you can create for a file system. The maximum number of SVMs depends on the amount of throughput capacity provisioned in megabytes per second (MBps).

Deployment type	Amount of throughput capacity (MBps)	Maximum number of SVMs per file system
Single-AZ (scale-up) and Multi-AZ (scale-up)	128	6
	256	6
	512	14
	1,024	14
	2,048	24
	4,096	24
Single-AZ (scale-out)	Any	5

Topics

- [Creating a storage virtual machine](#)
- [Updating a storage virtual machine](#)
- [Deleting a storage virtual machine \(SVM\)](#)
- [Viewing storage virtual machine configuration details](#)

Creating a storage virtual machine

You can create an FSx for ONTAP SVM using the Amazon Web Services Management Console, Amazon CLI, and API.

The maximum number of SVMs you can create for a file system depends on your file system's deployment type and the amount of throughput capacity provisioned. For more information, see [Maximum number of SVMs per file system](#).

SVM properties

When creating an SVM, you define the following properties:

- The FSx for ONTAP file system to which it belongs.

- The Microsoft Active Directory (AD) configuration – You can optionally join your SVM to a self-managed AD for authentication and access control of Windows and macOS clients. For more information, see [Working with Microsoft Active Directory in FSx for ONTAP](#).
- The root volume security style – Set the root volume security style (Unix, NTFS, or Mixed) to align with the type of clients that you're using to access your data within the SVM. For more information, see [Volume security style](#).
- The SVM administrative password – you can optionally set the password for the SVM's vsadmin user. For more information, see [Managing SVMs with the ONTAP CLI](#).

To create a storage virtual machine (console)

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. In the left navigation pane, choose **Storage virtual machines**.
3. Choose **Create new storage virtual machine**.

The **Create new storage virtual machine** dialog box appears.

Create new storage virtual machine ✕

File System

Select a filesystem ▼

Storage virtual machine name

Maximum of 47 alphanumeric characters, plus . - _ .

SVM administrative password
 Password for this SVM's "vsadmin" user, which you can use to access the ONTAP CLI or REST API.

Don't specify a password

Specify a password

Active Directory
 Joining an Active Directory enables access from Windows and MacOS clients over the SMB protocol.

Do not join an Active Directory

Join an Active Directory

Net BIOS name

Active Directory domain name
 This is the fully qualified domain name of your self-managed directory

example.com

DNS server IP addresses
 IPv4 addresses of the DNS servers for your domain

10.0.0.1

10.0.0.2 - optional

10.0.0.3 - optional

Service account username
 The username of the service account in your existing Active Directory. Do not include a domain prefix or suffix.

FSxServiceAccount

Service account password
 The password for the service account provided above.

Maximum of 128 characters.

Confirm password

Organizational Unit (OU) within which you want to join your file system - optional
 Specify the distinguished path name of the OU here

OU=org,DC=example,DC=com

Ensure that the service account provided has permissions delegated to the above OU or to the default OU if none is provided.

4. For **File system**, choose the file system to create the storage virtual machine on.
5. In the **Storage virtual machine name** field, provide a name for the storage virtual machine. You can use a maximum of 47 alphanumeric characters, plus the underscore (`_`) special character.
6. For **SVM administrative password**, you can optionally choose **Specify a password** and provide a password for this SVM's `vsadmin` user. You can use the `vsadmin` user to administer the SVM using the ONTAP CLI or REST API. For more information about the `vsadmin` user, see [Managing SVMs with the ONTAP CLI](#).

If you choose **Don't specify a password** (the default), you can still use the file system's `fsxadmin` user to manage your file system using the ONTAP CLI or REST API, but you can't use your SVM's `vsadmin` user to do the same.

7. For **Active Directory**, you have the following options:
 - If you are not joining your file system to an Active Directory (AD), choose **Do not join an Active Directory**.
 - If you are joining your SVM to a self-managed AD domain, choose **Join an Active Directory**, and provide the following details for your AD. For more information, see [Prerequisites for joining an SVM to a self-managed Microsoft AD](#).
 - The NetBIOS name of the Active Directory computer object to create for your SVM. The NetBIOS name cannot exceed 15 characters. This is the name of this SVM in Active Directory.
 - The fully qualified domain name (FQDN) of your Active Directory. The FQDN cannot exceed 255 characters.
 - **DNS server IP addresses** – The IPv4 addresses of the DNS servers for your domain.
 - **Service account username** – The username of the service account in your existing Active Directory. Do not include a domain prefix or suffix. For `EXAMPLE\ADMIN`, use `ADMIN`.
 - **Service account password** – The password for the service account.
 - **Confirm password** – The password for the service account.
 - (Optional) **Organizational Unit (OU)** – The distinguished path name of the organizational unit to which you want to join your file system.
 - **Delegated file system administrators group** – The name of the group in your AD that can administer your file system.

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

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

- For **SVM root volume security style**, choose the security style for the SVM depending on the type of clients that access your data. Choose **Unix (Linux)** if you primarily access your data using Linux clients; choose **NTFS** if you primarily access your data using Windows clients. For more information, see [Volume security style](#).
- Choose **Confirm** to create the storage virtual machine.

You can monitor the update progress on the **File systems** detail page, in the **Status** column of the **Storage virtual machines** pane. The storage virtual machine is ready for use when its status is **Created**.

To create a storage virtual machine (CLI)

- To create an FSx for ONTAP storage virtual machine (SVM), use the [create-storage-virtual-machine](#) CLI command (or the equivalent [CreateStorageVirtualMachine](#) API operation), as shown in the following example.

```
aws fsx create-storage-virtual-machine \
  --file-system-id fs-0123456789abcdef0 \
  --name svm1 \
  --svm-admin-password password \
  --active-directory-configuration
  SelfManagedActiveDirectoryConfiguration='{DomainName="corp.example.com", \
  OrganizationalUnitDistinguishedName="OU=FileSystems,DC=corp,DC=example,DC=com",FileSystemAd
  \
  UserName="FSxService",Password="password", \
  DnsIps=["10.0.1.18"]}',NetBiosName=amznfsx12345
```

After successfully creating the storage virtual machine, Amazon FSx returns its description in JSON format, as shown in the following example.

```
{
```

```

"StorageVirtualMachine": {
  "CreationTime": 1625066825.306,
  "Endpoints": {
    "Management": {
      "DnsName": "svm-abcdef0123456789a.fs-0123456789abcdef0.fsx.us-
east-1.amazonaws.com",
      "IpAddresses": ["198.19.0.4"]
    },
    "Nfs": {
      "DnsName": "svm-abcdef0123456789a.fs-0123456789abcdef0.fsx.us-
east-1.amazonaws.com",
      "IpAddresses": ["198.19.0.4"]
    },
    "Smb": {
      "DnsName": "amznfsx12345",
      "IpAddresses": ["198.19.0.4"]
    },
    "SmbWindowsInterVpc": {
      "IpAddresses": ["198.19.0.5", "198.19.0.6"]
    },
    "Iscsi": {
      "DnsName": "iscsi.svm-abcdef0123456789a.fs-0123456789abcdef0.fsx.us-
east-1.amazonaws.com",
      "IpAddresses": ["198.19.0.7", "198.19.0.8"]
    }
  },
  "FileSystemId": "fs-0123456789abcdef0",
  "Lifecycle": "CREATING",
  "Name": "vol1",
  "ResourceARN": "arn:aws:fsx:us-east-1:123456789012:storage-virtual-machine/
fs-0123456789abcdef0/svm-abcdef0123456789a",
  "StorageVirtualMachineId": "svm-abcdef0123456789a",
  "Subtype": "default",
  "Tags": [],
  "ActiveDirectoryConfiguration": {
    "NetBiosName": "amznfsx12345",
    "SelfManagedActiveDirectoryConfiguration": {
      "UserName": "Admin",
      "DnsIps": [
        "10.0.1.3",
        "10.0.91.97"
      ],
      "OrganizationalUnitDistinguishedName": "OU=Computers,OU=customer-
ad,DC=customer-ad,DC=example,DC=com",

```

```
        "DomainName": "customer-ad.example.com"
    }
}
}
```

Updating a storage virtual machine

You can update the following storage virtual machine (SVM) configuration properties using the Amazon FSx console, Amazon CLI, and Amazon FSx API:

- SVM administrative account password.
- SVM Active Directory (AD) configuration – You can join an SVM to an AD, or modify the AD configuration of an SVM already joined to an AD. For more information, see [Managing SVM Active Directory configurations](#).

To update the SVM administrator account credentials (console)

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. Choose the SVM to update as follows:
 - In the left navigation pane, choose **File systems**, and then choose the ONTAP file system for which you want to update an SVM.
 - Choose the **Storage virtual machines** tab.

–Or–

 - To display a list of all the SVMs available in your Amazon Web Services account in the current Amazon Web Services Region, expand **ONTAP** and choose **Storage virtual machines**.
3. Choose the storage virtual machine that you want to update.
4. Choose **Actions > Update administrator password**. The **Update SVM administrative credentials** window appears.
5. Enter the new password for the vsadmin user, and confirm it.
6. Choose **Update credentials** to save the new password.

To update the SVM administrator account credentials (CLI)

- To update the configuration of an FSx for ONTAP SVM, use the [update-storage-virtual-machine](#) CLI command (or the equivalent [UpdateStorageVirtualMachine](#) API operation), as shown in the following example.

```
aws fsx update-storage-virtual-machine \  
--storage-virtual-machine-id svm-abcdef01234567890 \  
--svm-admin-password new-svm-password \  

```

After successfully creating the storage virtual machine, Amazon FSx returns its description in JSON format, as shown in the following example.

```
{  
  "StorageVirtualMachine": {  
    "CreationTime": 1625066825.306,  
    "Endpoints": {  
      "Management": {  
        "DnsName": "svm-abcdef01234567890.fs-0123456789abcdef0.fsx.us-  
east-1.amazonaws.com",  
        "IpAddresses": ["198.19.0.4"]  
      },  
      "Nfs": {  
        "DnsName": "svm-abcdef01234567890.fs-0123456789abcdef0.fsx.us-  
east-1.amazonaws.com",  
        "IpAddresses": ["198.19.0.4"]  
      },  
      "Smb": {  
        "DnsName": "amznfsx12345",  
        "IpAddresses": ["198.19.0.4"]  
      },  
      "SmbWindowsInterVpc": {  
        "IpAddresses": ["198.19.0.5", "198.19.0.6"]  
      },  
      "Iscsi": {  
        "DnsName": "iscsi.svm-abcdef01234567890.fs-0123456789abcdef0.fsx.us-  
east-1.amazonaws.com",  
        "IpAddresses": ["198.19.0.7", "198.19.0.8"]  
      }  
    },  
    "FileSystemId": "fs-0123456789abcdef0",  
  }  
}
```

```
"Lifecycle": "CREATING",
  "Name": "vol1",
  "ResourceARN": "arn:aws:fsx:us-east-1:123456789012:storage-virtual-machine/
fs-0123456789abcdef0/svm-abcdef01234567890",
  "StorageVirtualMachineId": "svm-abcdef01234567890",
  "Subtype": "default",
  "Tags": [],
  "ActiveDirectoryConfiguration": {
    "NetBiosName": "amznfsx12345",
    "SelfManagedActiveDirectoryConfiguration": {
      "UserName": "Admin",
      "DnsIps": [
        "10.0.1.3",
        "10.0.91.97"
      ],
      "OrganizationalUnitDistinguishedName": "OU=Computers,OU=customer-
ad,DC=customer-ad,DC=example,DC=com",
      "DomainName": "customer-ad.example.com"
    }
  }
}
```

Deleting a storage virtual machine (SVM)

You can only delete an FSx for ONTAP SVM by using the Amazon FSx console, the Amazon CLI, and API. Before you can delete an SVM, you must delete all non-root volumes attached to the SVM first.

Important

You cannot delete an SVM by using the NetApp ONTAP CLI or API.

Note

Before you delete a storage virtual machine, make sure that no applications are accessing the data in the SVM, and that you have deleted all non-root volumes attached to the SVM.

To delete a storage virtual machine (console)

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. Choose the SVM that you want to delete as follows:
 - In the left navigation pane, choose **File systems**, and then choose the ONTAP file system for which you want to delete an SVM.
 - Choose the **Storage virtual machines** tab.

–Or–

- To display a list of all the SVMs available, expand **ONTAP** and choose **Storage virtual machines**.

Select the SVM that you want to delete from the list.

3. In the **Volumes** tab, view the list of volumes attached to the SVM. If there are any non-root volumes attached to the SVM, you must delete them before you can delete the SVM. See [Deleting a volume](#) for more information.
4. Choose **Delete storage virtual machine** from the **Actions** menu.
5. In the delete confirmation dialog box, choose **Delete storage virtual machine**.

To delete a storage virtual machine (CLI)

- To delete an FSx for ONTAP storage virtual machine, use the [delete-storage-virtual-machine](#) CLI command (or the equivalent [DeleteStorageVirtualMachine](#) API operation), as shown in the following example.

```
aws fsx delete-storage-virtual-machine --storage-virtual-machine-id svm-  
abcdef0123456789d
```

Viewing storage virtual machine configuration details

You can see the FSx for ONTAP storage virtual machines that are currently on your file system using the Amazon FSx console, the Amazon CLI, and the Amazon FSx API.

To view a storage virtual machine on your file system:

- **Using the console** – Choose a file system to view its **File systems** detail page. To list all the storage virtual machines on the file system, choose the **Storage virtual machines** tab, and then choose the storage virtual machine that you want to view.
- **Using the CLI or API** – Use the [describe-storage-virtual-machines](#) CLI command or the [DescribeStorageVirtualMachines](#) API operation.

The system response is a list of full descriptions of all the SVMs in your account in that Amazon Web Services Region.

Managing FSx for ONTAP volumes

Each storage virtual machine (SVM) on an FSx for ONTAP file system can have one or more *volumes*. A volume is an isolated data container for files, directories, or iSCSI logical units of storage (LUNs). Volumes are *thin provisioned*, meaning that they consume storage capacity only for the data stored in them.

You can access a volume from Linux, Windows, or macOS clients over the Network File System (NFS) protocol, the Server Message Block (SMB) protocol, or over the Internet Small Computer Systems Interface (iSCSI) protocol by creating an iSCSI LUN (shared block storage). FSx for ONTAP also supports multi-protocol access (concurrent NFS and SMB access) to the same volume.

You can create volumes by using the Amazon Web Services Management Console, Amazon CLI, the Amazon FSx API, or NetApp BlueXP. You can also use your file system's or SVM's administrative endpoint to create, update, and delete volumes by using the NetApp ONTAP CLI or REST API.

You can create up to 500 volumes per scale-up file system and 1,000 volumes per scale-out file system.

When you create a volume, you define the following properties:

- **Volume style** – The [volume style](#) can be either FlexVol or FlexGroup.
- **Volume name** – The name of the volume.
- **Volume type** – The [volume type](#) can be either Read-Write (RW) or Data protection (DP). DP volumes are read-only and used as the destination in a NetApp SnapMirror or SnapVault relationship.

- Volume size – This is the maximum amount of data that the volume can store, regardless of the storage tier.
- Junction path – This is the location in the SVM's namespace where the volume gets mounted.
- Storage efficiency – [Storage efficiency](#) features, including data compaction, compression, and deduplication provide typical storage savings of 65% for general-purpose file sharing workloads.
- Volume [security style](#) (Unix, NTFS, or Mixed) – Determines what type of permissions are used for data access on the volume when authorizing users.
- Data tiering – The [tiering policy](#) defines which data is stored in the cost-effective capacity pool tier.
- [Tiering policy cooling period](#) – Defines when data is marked cold and moved to capacity pool storage.
- Snapshot policy – [Snapshot policies](#) define how the system creates snapshots for a volume. You can choose from three predefined policies or use a custom policy that you have created using the ONTAP CLI or REST API.
- [Copy tags to backups](#) – Amazon FSx will automatically copy any tags from your volumes to backups using this option. You can set this option using the Amazon CLI or Amazon FSx API.

Topics

- [Volume styles](#)
- [Volume types](#)
- [Volume security style](#)
- [Creating volumes](#)
- [Updating a volume](#)
- [Deleting a volume](#)
- [Viewing a volume](#)

Volume styles

FSx for ONTAP offers two styles of volumes that you can use for different purposes. You can create either FlexVol or FlexGroup volumes using the Amazon FSx console, the Amazon CLI, and the Amazon FSx API.

- FlexVol volumes offer the simplest experience for file systems with one high-availability (HA) pair and are the default volume style for scale-up file systems. The minimum size of a FlexVol volume is 20 mebibytes (MiB), and the maximum size is 314,572,800 MiB.
- FlexGroup volumes are comprised of multiple constituent FlexVol volumes, which allows them to deliver higher performance and storage scalability than FlexVol volumes for file systems with multiple HA pairs. FlexGroup volumes are the default volume style for scale-out file systems. The minimum size of a FlexGroup volume is 100 gibibytes (GiB) per constituent, and the maximum size is 20 pebibytes (PiB).

You can convert a volume with the FlexVol style to the FlexGroup style with the ONTAP CLI, which creates a FlexGroup with a single constituent. However, we recommend that you use Amazon DataSync to move data between a FlexVol volume and a new FlexGroup volume to ensure that the data is evenly distributed across the FlexGroup's constituents. For more information, see [FlexGroup constituents](#).

Note

If you want to use the ONTAP CLI to convert a FlexVol volume to a FlexGroup volume, make sure that you delete any backups of the FlexVol volume before converting it. ONTAP doesn't automatically rebalance data as part of the conversion, so the data might be imbalanced across the FlexGroup constituents.

FlexGroup constituents

A FlexGroup volume is made up of constituents, which are FlexVol volumes. By default, FSx for ONTAP assigns eight constituents to a FlexGroup volume per HA pair.

When you create your FlexGroup volume, the size of it is divided evenly among its constituents. For example, if you create an 800 gigabyte (GB) FlexGroup volume with eight constituents, each constituent is 100 GB in size. A FlexGroup volume can be between 100 GB and 20 PiB in size, but the total size depends on the size of the constituents. Each constituent has a minimum size of 100 GB and a maximum size of 300 TiB. For example, a FlexGroup volume with eight constituents has a minimum size of 800 GB and a maximum size of 20 PiB.

ONTAP distributes data at the file-level across the constituents. You can store up to two billion files in each constituent on your FlexGroup volume.

When you update the size of your FlexGroup volume, the new size is evenly distributed among its existing constituents.

You can also add more constituents to your FlexGroup volume using the ONTAP CLI or REST API. However, we recommend that you only do so if you need additional storage capacity and all of your constituents are already at their maximum size (300 TiB per constituent). Adding constituents can lead to an imbalance of data and I/O across the constituents. Until the constituents are balanced, it's possible that the write throughput might be 5–10% lower than a balanced FlexGroup volume. When new data is written to the FlexGroup volume, ONTAP prioritizes distributing it among the new constituents until the constituents are balanced. If you do add new constituents, we recommend choosing an even number and not exceeding eight per aggregate.

Note

If you add new constituents, your existing snapshots become partial snapshots; therefore, they can't be used to fully restore your FlexGroup volume to a prior state. The previous snapshots can't offer a complete point-in-time image of your FlexGroup volume because the new constituents didn't exist yet. However, the partial snapshots can be used to restore individual files and directories, to create a new volume, or to replicate with SnapMirror.

Volume types

FSx for ONTAP offers two types of volumes that you can create using the Amazon FSx console, the Amazon CLI, and the Amazon FSx API.

- Read-write (RW) volumes are used in most cases. As their name indicates, they are read-writable.
- Data protection (DP) volumes are read-only volumes that you use as the destination of a NetApp SnapMirror or SnapVault relationship. You should use DP volumes when you want to [migrate](#) or [protect](#) a single volume's data.

FlexVol and FlexGroup volumes can be either RW or DP.

Note

You can't update a volume's type after the volume is created.

Volume security style

FSx for ONTAP supports 3 different volume security styles: Unix, NTFS, and mixed. Each security style has a different effect on how permissions are handled for data. You must understand the different effects to ensure that you select the appropriate security style for your purposes.

It is important to understand that security styles do not determine what client types can or cannot access data. Security styles only determine the type of permissions FSx for ONTAP uses to control data access and what client type can modify these permissions.

The two factors that you use to determine the security style for a volume are the type of administrators that manage the file system and the type of users or services that access the data on the volume.

When creating a volume in the Amazon FSx console, CLI, and API, the security style is automatically set to the root volume's security style. You can modify a volume's security style using the Amazon CLI or API. You can modify this setting after the volume is created. See [Updating a volume](#) for more information.

When you configure the security style on a volume, consider the needs of your environment to ensure that you select the best security style in order to avoid issues with managing permissions. Keep in mind that security style doesn't determine which client types can access data. Security style determines the permissions that are used to allow data access and the client types that can modify those permissions. Following are considerations that can help you decide which security style to choose for a volume:

- **Unix (Linux)** – Choose this security style if the file system is managed by a Unix administrator, the majority of users are NFS clients, and an application accessing the data uses a Unix user as the service account. Only Linux clients can modify permissions with the Unix security style, and the type of permissions used on files and directories are mode-bits or NFS v4.x ACLs.
- **NTFS** – Choose this security style if the file system is managed by a Windows administrator, the majority of users are SMB clients, and an application accessing the data uses a Windows user as the service account. If any Windows access is required to a volume, we recommend that you use the NTFS security style. Only Windows clients can modify permissions with NTFS security style, and the types of permissions used on file and directories is NTFS ACLs.
- **Mixed** – This is an advanced setting. For more information, see the topic [What the security styles and their effects are](#) in the NetApp Documentation Center.

Creating volumes

You can create an FSx for ONTAP FlexVol or FlexGroup volume using the Amazon FSx console, the Amazon CLI, and the Amazon FSx API, in addition to the NetApp ONTAP command line interface (CLI) and REST API.

To create a FlexVol volume (console)

Note

The volume's security style is automatically set to the root volume's security style.

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. In the left navigation pane, choose **Volumes**.
3. Choose **Create volume**.
4. For **File system type**, choose **Amazon FSx for NetApp ONTAP**.
5. In the **File system details** section, provide the following information:
 - For **File system**, choose the file system to create the volume on.
 - For **Storage virtual machine**, choose the storage virtual machine (SVM) to create the volume on.
6. In the **Volume style** section, choose **FlexVol**.
7. In the **Volume details** section, provide the following information:
 - In the **Volume name** field, provide a name for the volume. You can use up to 203 alphanumeric or underscore (_) characters.
 - For **Volume size**, enter any whole number in the range of 20–314572800 to specify the size in mebibytes (MiB).
 - For **Volume type**, choose **Read-Write (RW)** to create a volume that is readable and writable or **Data Protection (DP)** to create a volume that is read-only and can be used as the destination of a NetApp SnapMirror or SnapVault relationship. For more information, see [Volume types](#).
 - For **Junction path**, enter a location within the file system to mount the volume. The name must have a leading forward slash, for example `/vo13`.

- For **Storage efficiency**, choose **Enabled** to enable the ONTAP storage-efficiency features (deduplication, compression, and compaction). For more information, see [FSx for ONTAP storage efficiency](#).
- For **Volume security style**, choose between **Unix (Linux)**, **NTFS**, and **Mixed** for the volume. For more information, see [Volume security style](#).
- For **Snapshot policy**, choose a snapshot policy for the volume. For more information about snapshot policies, see [Snapshot policies](#).

If you choose **Custom policy**, you must specify the policy's name in the **custom-policy** field. The custom policy must already exist on the SVM or in the file system. You can create a custom snapshot policy with the ONTAP CLI or REST API. For more information, see [Create a Snapshot Policy](#) in the NetApp ONTAP Product Documentation.

8. In the **Storage tiering** section, provide the following information:
 - For **Capacity pool tiering policy**, choose the storage pool tiering policy for the volume, which can be **Auto** (the default), **Snapshot Only**, **All**, or **None**. For more information, see [Volume tiering policies](#).
 - If you choose either **Auto** or **Snapshot Only**, you can set the **Tiering policy cooling period** to define the number of days before data that has not been accessed is marked cold and moved to capacity pool storage. You can provide a value between 2 and 183 days. The default setting is 31 days.
9. In the **Advanced** section, for **SnapLock Configuration**, choose between **Enabled** and **Disabled**. For more information about configuring a SnapLock Compliance volume or a SnapLock Enterprise volume, see [Creating a SnapLock Compliance volume](#) and [Creating a SnapLock Enterprise volume](#). For more information about SnapLock, see [Protecting your data with SnapLock](#).
10. Choose **Confirm** to create the volume.

You can monitor the update progress on the **File systems** detail page, in the **Status** column of the **Volumes** pane. The volume is ready for use when its status is **Created**.


To create a FlexGroup volume (console)

Note

You can only create FlexGroup volumes for scale-out file systems using the Amazon FSx console. To create FlexVol volumes for your scale-out file systems, use the Amazon CLI, Amazon FSx API, or NetApp management tools.

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. In the left navigation pane, choose **Volumes**.
3. Choose **Create volume**.
4. For **File system type**, choose **Amazon FSx for NetApp ONTAP**.
5. In the **File system details** section, provide the following information:
 - For **File system**, choose the file system to create the volume on.
 - For **Storage virtual machine**, choose the storage virtual machine (SVM) to create the volume on.
6. In the **Volume style** section, choose **FlexGroup**.
7. In the **Volume details** section, provide the following information:
 - In the **Volume name** field, provide a name for the volume. You can use up to 203 alphanumeric or underscore (_) characters.
 - For **Volume size**, enter any whole number in the range of 800 gibibytes (GiB)–2,000 pebibytes (PiB).
 - For **Volume type**, choose **Read-Write (RW)** to create a volume that is readable and writable or **Data Protection (DP)** to create a volume that is read-only and can be used as the destination of a NetApp SnapMirror or SnapVault relationship. For more information, see [Volume types](#).
 - For **Junction path**, enter a location within the file system to mount the volume. The name must have a leading forward slash, for example /vo13.
 - For **Storage efficiency**, choose **Enabled** to enable the ONTAP storage-efficiency features (deduplication, compression, and compaction). For more information, see [FSx for ONTAP storage efficiency](#).

- For **Volume security style**, choose between **Unix (Linux)**, **NTFS**, and **Mixed** for the volume. For more information, see [Volume security style](#).

 **Note**

The volume's security style is automatically set to the root volume's security style.

- For **Snapshot policy**, choose a snapshot policy for the volume. For more information about snapshot policies, see [Snapshot policies](#).

If you choose **Custom policy**, you must specify the policy's name in the **custom-policy** field. The custom policy must already exist on the SVM or in the file system. You can create a custom snapshot policy with the ONTAP CLI or REST API. For more information, see [Create a Snapshot Policy](#) in the NetApp ONTAP Product Documentation.

8. In the **Storage tiering** section, provide the following information:
 - For **Capacity pool tiering policy**, choose the storage pool tiering policy for the volume, which can be **Auto** (the default), **Snapshot Only**, **All**, or **None**. For more information, see [Volume tiering policies](#).
 - If you choose either **Auto** or **Snapshot Only**, you can set the **Tiering policy cooling period** to define the number of days before data that has not been accessed is marked cold and moved to capacity pool storage. You can provide a value between 2–183 days. The default setting is 31 days.
9. In the **Advanced** section, for **SnapLock Configuration**, choose between **Enabled** and **Disabled**. For more information about configuring a SnapLock Compliance volume or a SnapLock Enterprise volume, see [Creating a SnapLock Compliance volume](#) and [Creating a SnapLock Enterprise volume](#). For more information about SnapLock, see [Protecting your data with SnapLock](#).
10. Choose **Confirm** to create the volume.

You can monitor the update progress on the **File systems** detail page, in the **Status** column of the **Volumes** pane. The volume is ready for use when its status is **Created**.

To create a volume (CLI)

- To create an FSx for ONTAP volume, use the [create-volume](#) CLI command (or the equivalent [CreateVolume](#) API operation), as shown in the following example.

```
aws fsx create-volume \
  --volume-type ONTAP \
  --name vol1 \
  --ontap-configuration CopyTagsToBackups=true,JunctionPath=/
vol1,SecurityStyle=NTFS, \
  SizeInMegabytes=1024,SnapshotPolicy=default, \
  StorageVirtualMachineId=svm-abcdef0123456789a,OntapVolumeType=RW, \
  StorageEfficiencyEnabled=true
```

After successfully creating the volume, Amazon FSx returns its description in JSON format, as shown in the following example.

```
{
  "Volume": {
    "CreationTime": "2022-08-12T13:03:37.625000-04:00",
    "FileSystemId": "fs-abcdef0123456789c",
    "Lifecycle": "CREATING",
    "Name": "vol1",
    "OntapConfiguration": {
      "CopyTagsToBackups": true,
      "FlexCacheEndpointType": "NONE",
      "JunctionPath": "/vol1",
      "SecurityStyle": "NTFS",
      "SizeInMegabytes": 1024,
      "SnapshotPolicy": "default",
      "StorageEfficiencyEnabled": true,
      "StorageVirtualMachineId": "svm-abcdef0123456789a",
      "StorageVirtualMachineRoot": false,
      "TieringPolicy": {
        "Name": "NONE"
      },
      "OntapVolumeType": "RW"
    },
    "ResourceARN": "arn:aws:fsx:us-east-2:111122223333:volume/fs-abcdef0123456789c/
fsvol-abcdef0123456789b",
    "VolumeId": "fsvol-abcdef0123456789b",
    "VolumeType": "ONTAP"
  }
}
```

```
}
```

You can also create a new volume by restoring a backup of a volume to a new volume. For more information, see [Restoring backups to a new volume](#).

Updating a volume

You can update the configuration of an FSx for ONTAP volume using the Amazon FSx console, the Amazon CLI, and the Amazon FSx API, in addition to the NetApp ONTAP command line interface (CLI) and REST API. You can modify the following properties of an existing FSx for ONTAP volume:

- Volume name
- Junction path
- Volume size
- Storage efficiency
- Capacity pool tiering policy
- Volume security style
- Snapshot policy
- Tiering policy cooling period
- Copy tags to backups (using the Amazon CLI and Amazon FSx API)

For more information, see [Managing FSx for ONTAP volumes](#).

To update a volume configuration (console)

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. Navigate to **File systems** and choose the ONTAP file system that you want to update a volume for.
3. Choose the **Volumes** tab.
4. Choose the volume that you want to update.
5. For **Actions**, choose **Update volume**.

The **Update volume** dialog box displays with the volume's current settings.

6. For **Junction path**, enter an existing location within the file system to mount the volume. The name must have a leading forward slash, such as `/vo15`.

7. For **Volume size**, you can increase or decrease the size of the volume within the range specified in the Amazon FSx console. For FlexVol volumes, the maximum size is 300 TiB. For FlexGroup volumes, the maximum size is 300 TiB multiplied by the total number of constituent volumes that your FlexGroup has, up to a maximum of 20 PiB.
8. For **Storage efficiency**, choose **Enabled** to enable the ONTAP storage efficiency features (deduplication, compression, and compaction), or choose **Disabled** to disable them.
9. For **Capacity pool tiering policy**, choose a new storage pool tiering policy for the volume, which can be **Auto** (the default), **Snapshot-only**, **All**, or **None**. For more information about capacity pool tiering policies, see [Volume tiering policies](#).
10. For **Volume security style**, choose either **Unix (Linux)**, **NTFS**, or **Mixed**. A volume's security style determines whether preference is given to NTFS or UNIX ACLs for multi-protocol access. The MIXED mode is not required for multi-protocol access and is only recommended for advanced users.
11. For **Snapshot policy**, choose a snapshot policy for the volume. For more information about snapshot policies, see [Snapshot policies](#).

If you choose **Custom policy**, you must specify the policy's name in the **custom-policy** field. The custom policy must already exist on the SVM or in the file system. You can create a custom snapshot policy with the ONTAP CLI or REST API. For more information, see [Create a Snapshot Policy](#) in the NetApp ONTAP Product Documentation.

12. For **Tiering policy cooling period**, valid values are 2-183 days. A volume's tiering policy cooling period defines the number of days before data that has not been accessed is marked cold and moved to capacity pool storage. This setting only affects the Auto and Snapshot-only policies.
13. Choose **Update** to update the volume.

To update a volume's configuration (CLI)

- To update the configuration of an FSx for ONTAP volume, use the [update-volume](#) CLI command (or the equivalent [UpdateVolume](#) API operation), as shown in the following example.

```
aws fsx update-volume \  
  --volume-id fsvol-1234567890abcdefa \  
  --name new_vol \  
  --ontap-configuration CopyTagsToBackups=true,JunctionPath=/new_vol, \  
  --
```

```
SizeInMegabytes=2048,SnapshotPolicy=default-1weekly, \  
StorageEfficiencyEnabled=true, \  
TieringPolicy=all
```

Deleting a volume

You can delete an FSx for ONTAP volume using the Amazon FSx console, the Amazon CLI, and the Amazon FSx API, in addition to the NetApp ONTAP command line interface (CLI) and REST API.

Important

You can only delete volumes using the Amazon FSx console, API, or CLI if the volume has Amazon FSx backups enabled.

Important

When you delete a volume by using the Amazon FSx console, you have the option to take a final backup of the volume. You can create new volumes from backups. We recommend that you choose to take a final backup as a best practice. If you find you don't need it after a certain period of time, you can delete this and other manually created volume backups. When you delete a volume by using the `delete-volume` CLI command, Amazon FSx takes a final backup by default.

Before you delete a volume, make sure that no applications are accessing the data in the volume that you want to delete.

To delete a volume (console)

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. In the left navigation pane, choose **File systems**, and then choose the ONTAP file system that you want to delete a volume from.
3. Choose the **Volumes** tab.
4. Choose the volume that you want to delete.
5. For **Actions**, choose **Delete volume**.
6. In the confirmation dialog box, for **Create final backup**, you have two options:

- Choose **Yes** to take a final backup of the volume. The name of the final backup is displayed.
 - Choose **No** if you don't want a final backup of the volume. You are asked to acknowledge that once the volume is deleted, automatic backups are no longer available.
7. Confirm the volume deletion by entering **delete** in the **Confirm delete** field.
 8. Choose **Delete volume(s)**.

To delete a volume (CLI)

- To delete an FSx for ONTAP volume, use the [delete-volume](#) CLI command (or the equivalent [DeleteVolume](#) API operation), as shown in the following example.

```
aws fsx delete-volume --volume-id fsvol-1234567890abcde
```

Viewing a volume

You can see the FSx for ONTAP volumes that are currently on your file system using the Amazon FSx console, the Amazon CLI, and the Amazon FSx API and SDKs.

To view the volumes on your file system:

- **Using the console** – Choose a file system to view the **File systems** detail page. Choose the **Volumes** tab to list all the volumes on the file system, and then choose the volume you want to view.
- **Using the CLI or API** – Use the [describe-volumes](#) CLI command or the [DescribeVolumes](#) API operation.

Creating an iSCSI LUN

This process describes how to create an iSCSI LUN on an Amazon FSx for NetApp ONTAP scale-up file system using the NetApp ONTAP CLI **lun create** command. For more information, see [lun create](#) in the NetApp ONTAP Documentation Center.

Note

The iSCSI protocol isn't supported for scale-out file systems.


This process assumes you already have a volume created on your file system. For more information, see [Creating volumes](#).

1. To access the NetApp ONTAP CLI, establish an SSH session on the management port of the Amazon FSx for NetApp ONTAP file system by running the following command. Replace *management_endpoint_ip* with the IP address of the file system's management port.

```
[~]$ ssh fsxadmin@management_endpoint_ip
```


For more information, see [Managing file systems with the ONTAP CLI](#).

2. Create a LUN using the **lun create** NetApp CLI command, replacing the following values:
 - **svm_name** - The name of the storage virtual machine (SVM) providing the iSCSI target. The host uses this value to reach the LUN.
 - **vol_name** - The name of the volume hosting the LUN.
 - **lun_name** - The name you want to assign to the LUN.
 - **size** - The size, in bytes, of the LUN. The maximum size LUN you can create is 128 TB.

 **Note**

We recommend that you use a volume at least 5% larger than your LUN size. This margin leaves space for volume snapshots.

- **ostype** - The operating system of the host, either `windows_2008` or `linux`. Use `windows_2008` for all versions of Windows; this ensures the LUN has proper block offset for the operating system and optimizes performance.

 **Note**

We recommend enabling space allocation on your LUN. With space allocation enabled, ONTAP can inform your host when the LUN is out of capacity and can reclaim space as you delete data from the LUN.

For more information, see [lun create](#) in the NetApp ONTAP CLI documentation.

```
> lun create -vserver svm_name -path /vol/vol_name/lun_name -size size -  
ostype ostype -space-allocation enabled
```

```
Created a LUN of size 10g (10737418240)
```

3. Confirm the LUN is created, online, and mapped.

```
> lun show
```

The system responds with the following output:

Vserver	Path	State	Mapped	Type	Size
<i>svm_name</i>	<i>/vol/vol_name/lun_name</i>	online	unmapped	windows_2008	10GB

Next steps

Now that you have created an iSCSI LUN, the next step in the process of using an iSCSI LUN as block storage is to map the LUN to an igroup. For more information, see [Mounting iSCSI LUNs to a Linux client](#) or [Mounting iSCSI LUNs to a Windows client](#).

Managing SMB shares

To manage SMB file shares on your Amazon FSx file system, you can use the Microsoft Windows Shared Folders GUI. The Shared Folders GUI provides a central location for managing all shared folders in your storage virtual machine (SVM). The following procedures detail how to create, update, and remove your file shares.

Note

You can also manage SMB file shares by using the NetApp System Manager. For more information, see [Using NetApp System Manager with BlueXP](#).

To connect shared folders to your Amazon FSx file system

1. Launch your Amazon EC2 instance and connect it to the Microsoft Active Directory that your Amazon FSx file system is joined to. To do this, choose one of the following procedures from the *Amazon Directory Service Administration Guide*:
 - [Seamlessly join a Windows EC2 instance](#)
 - [Manually join a Windows instance](#)
2. Connect to your instance as a user that is a member of the file system administrators group. For more information, see [Connecting to Your Windows Instance](#) in the *Amazon EC2 User Guide for Windows Instances*.
3. Open the **Start** menu and run **fsmgmt.msc** using **Run As Administrator**. Doing this opens the Shared Folders GUI tool.
4. For **Action**, choose **Connect to another computer**.
5. For **Another computer**, enter the DNS name for your storage virtual machine (SVM), for example **netbios_name.corp.example.com**.

To find your SVM's DNS name on the Amazon FSx console, choose **Storage virtual machines**, choose your SVM, and then scroll down to **Endpoints** until you find **SMB DNS name**. You can also get the DNS name in the response of the [DescribeStorageVirtualMachines](#) API operation.

6. Choose **OK**. An entry for your Amazon FSx file system then appears in the list for the Shared Folders tool.

Now that Shared Folders is connected to your Amazon FSx file system, you can manage the Windows file shares on the file system with the following actions:

Note

We recommend that you locate your SMB shares on a volume other than your root volume.

- **Create a new file share** – In the Shared Folders tool, choose **Shares** in the left pane to see the active shares for your Amazon FSx file system. Volumes are shown mounted on the path chosen during volume creation. Choose **New Share** and complete the Create a Shared Folder wizard.

You have to create the local folder prior to creating the new file share. You can do so as follows:

- Using the Shared Folders tool: choose **Browse** when specifying a local folder path, choose **Make new folder** to create the local folder.
- Using command line:

```
New-Item -Type Directory -Path \\netbios_name.corp.example.com\C
$volume_path\MyNewFolder
```

- **Modify a file share** – In the Shared Folders tool, open the context (right-click) menu for the file share that you want to modify in the right pane, and choose **Properties**. Modify the properties and choose **OK**.
- **Remove a file share** – In the Shared Folders tool, open the context (right-click) menu for the file share that you want to remove in the right pane, and then choose **Stop Sharing**.

Note

Removing file shares from the GUI is possible only if you connected to **fsmgmt.msc** using the DNS name of the Amazon FSx file system. If you connected using the IP address or DNS alias name of the file system, the **Stop Sharing** option won't work and the file share isn't removed.

File access auditing

Amazon FSx for NetApp ONTAP supports auditing of end-user accesses to files and directories in a storage virtual machine (SVM).

Topics

- [File access auditing overview](#)
- [Overview of tasks for setting up file access auditing](#)

File access auditing overview

File access auditing enables you to record end-user accesses of individual files and directories based on audit policies you define. File access auditing can help you improve your system's security and reduce the risk of unauthorized access to your system data. File access auditing helps your organizations remain compliant with data protection requirements, identify potential threats early, and reduce the risk of a data breach.


Across file and directory accesses, Amazon FSx supports logging of successful attempts (such as a user with sufficient permissions successfully accessing a file), failed attempts, or both. You can also turn off file access auditing at any time.

By default, audit event logs are stored in the EVTX file format, which allows you to view them using Microsoft Event Viewer.

SMB access events that can be audited

The following table lists the SMB file and folder access events can be audited.

Event ID (EVT/EVTX)	Event	Description	Category
560/4656	Open Object/Create Object	OBJECT ACCESS: Object (file or directory) open	File Access
563/4659	Open Object with the Intent to Delete	OBJECT ACCESS: A handle to an object (file or directory) was requested with the Intent to Delete	File Access
564/4660	Delete Object	OBJECT ACCESS: Delete Object (file or directory). ONTAP generates this event when a Windows client attempts to delete the object (file or directory)	File Access
567/4663	Read Object/Write Object/Get Object Attributes/Set Object Attributes	OBJECT ACCESS: Object access attempt (read, write, get attribute, set attribute).	File Access

Event ID (EVT/EVTX)	Event	Description	Category
		<p> Note</p> <p>For this event, ONTAP audits only the first SMB read and first SMB write operation (success or failure) on an object. This prevents ONTAP from creating excessive log entries when a single client opens an object and performs many successive read or write operations to the same object.</p>	
N/A/4664	Hard link	OBJECT ACCESS: An attempt was made to create a hard link	File Access

Event ID (EVT/EVTX)	Event	Description	Category
N/A/N/A ONTAP Event ID 9999	Rename Object	OBJECT ACCESS: Object renamed. This is an ONTAP event. It is not currently supported by Windows as a single event.	File Access
N/A/N/A ONTAP Event ID 9998	Unlink Object	OBJECT ACCESS: Object unlinked. This is an ONTAP event. It is not currently supported by Windows as a single event.	File Access

NFS access events that can be audited

The following NFS file and folder access events can be audited.

- READ
- OPEN
- CLOSE
- READDIR
- WRITE
- SETATTR
- CREATE
- LINK
- OPENATTR
- REMOVE
- GETATTR
- VERIFY

- NVERIFY
- RENAME

Overview of tasks for setting up file access auditing

Setting up FSx for ONTAP for file access auditing involves the following high-level tasks:

1. [Familiarize yourself](#) with the file access auditing requirements and considerations.
2. [Create an auditing configuration](#) on a specific SVM.
3. [Enable auditing](#) on that SVM.
4. [Configure audit policies](#) on your files and directories.
5. [View the audit event logs](#) after FSx for ONTAP emits them.

Task details are provided in the following procedures.

Repeat the tasks for any other SVM on your file system that you want to enable file access auditing for.

Auditing requirements

Before you configure and enable auditing on an SVM, you should be aware of the following requirements and considerations.

- NFS auditing supports audit Access Control Entries (ACEs) designated as type u, which generate an audit log entry when access is attempted on the object. For NFS auditing, there is no mapping between mode bits and audit ACEs. When converting ACLs to mode bits, audit ACEs are skipped. When converting mode bits to ACLs, audit ACEs are not generated.
- Auditing is dependent on having available space in the staging volumes. (A staging volume is dedicated volume created by ONTAP to store staging files, which are intermediate binary files on individual nodes where audit records are stored prior to conversion to an EVTX or XML file format.) You must ensure that there is sufficient space for the staging volumes in aggregates that contain audited volumes.
- Auditing is dependent on having available space in the volume containing the directory where converted audit event logs are stored. You must ensure that there is sufficient space in the volumes used to store event logs. You can specify the number of audit logs to retain

in the auditing directory by using the `-rotate-limit` parameter when creating an auditing configuration, which can help to ensure that there is enough available space for the audit logs in the volume.

Creating auditing configurations on SVMs

Before you can begin auditing file and directory events, you must create an auditing configuration on the Storage Virtual Machine (SVM). After you create the auditing configuration, you must enable it on the SVM.

Before you use the `vserver audit create` command to create the auditing configuration, make sure you've created a directory to be used as the destination for logs, and that the directory doesn't have symlinks. You specify the destination directory with the `-destination` parameter.

You can create an auditing configuration that rotates audit logs based on log size or a schedule, as follows:

- To rotate audit logs based on log size, use this command:

```
vserver audit create -vserver svm_name -destination path [-format {xml|evtx}] [-rotate-limit integer] [-rotate-size {integer[KB|MB|GB|TB|PB]}]
```

The following example creates an auditing configuration for the SVM named `svm1` that audits file operations and CIFS (SMB) logon and logoff events (the default) using size-based rotation. The log format is EVTX (the default), logs are stored in the `/audit_log` directory, and you'll have a single log file at a time (up to 200MB in size).

```
vserver audit create -vserver svm1 -destination /audit_log -rotate-size 200MB
```

- To rotate audit logs based on a schedule, use this command:

```
vserver audit create -vserver svm_name -destination path [-format {xml|evtx}]  
  [-rotate-limit integer] [-rotate-schedule-month chron_month]  
  [-rotate-schedule-dayofweek chron_dayofweek] [-rotate-schedule-  
day chron_dayofmonth]  
  [-rotate-schedule-hour chron_hour] [-rotate-schedule-minute chron_minute]
```

The `-rotate-schedule-minute` parameter is required if you are configuring time-based audit log rotation.

The following example creates an auditing configuration for the SVM named `svm2` using time-based rotation. The log format is EVTX (the default) and the audit logs are rotated monthly, at 12:30 PM on all days of the week.

```
vserver audit create -vserver svm2 -destination /audit_log -rotate-size 200MB -  
rotate-schedule-month all -rotate-schedule-dayofweek all -rotate-schedule-hour 12 -  
rotate-schedule-minute 30
```

You can use the `-format` parameter to specify whether the audit logs are created in the converted EVTX format (the default) or in the XML file format. The EVTX format allows you to view the log files with Microsoft Event Viewer.

By default, the categories of events to be audited are file access events (both SMB and NFS), CIFS (SMB) logon and logoff events, and authorization policy change events. You can have greater control over which events to log by the `-events` parameter, which has the following format:

```
-events {file-ops|cifs-logon-logoff|cap-staging|file-share|audit-policy-change|user-  
account|authorization-policy-change|security-group}
```

For example, using `-events file-share` enables auditing of file share events.

For more information on the `vserver audit create` command, see [Create an audit configuration](#).

Enabling auditing on an SVM

After you finish setting up the auditing configuration, you must enable auditing on the SVM. To do so, use the following command:

```
vserver audit enable -vserver svm_name
```

For example, use the following command to enable auditing on the SVM named `svm1`.

```
vserver audit enable -vserver svm1
```

You can disable access auditing at any time. For example, use the following command to turn off auditing on the SVM named `svm4`.

```
vserver audit disable -vserver svm4
```

When you disable auditing, the audit configuration isn't deleted on the SVM, which means that you can re-enable auditing on that SVM at any time.

Configuring file and folder audit policies

You need to configure audit policies on the files and folders that you want audited for user access attempts. You can configure audit policies to monitor both successful and failed access attempts.

You can configure both SMB and NFS audit policies. SMB and NFS audit policies have different configuration requirements and audit capabilities based on the security style of the volume.

Audit policies on NTFS security-style files and directories

You can configure NTFS audit policies by using the Windows Security tab or the ONTAP CLI.

To configure NTFS audit policies (Windows Security tab)

You configure NTFS audit policies by adding entries to NTFS SACLs that are associated with an NTFS security descriptor. The security descriptor is then applied to NTFS files and directories. These tasks are automatically handled by the Windows GUI. The security descriptor can contain discretionary access control lists (DACLS) for applying file and folder access permissions, SACLs for file and folder auditing, or both SACLs and DACLS.

1. From the **Tools** menu in Windows Explorer, select **Map network drive**.
2. Complete the **Map Network Drive** box:
 - a. Choose a **Drive** letter.
 - b. In the **Folder** box, type the SMB (CIFS) server name that contains the share, holding the data you want to audit and the name of the share.
 - c. Choose **Finish**.

The drive you selected is mounted and ready with the Windows Explorer window displaying files and folders contained within the share.

3. Select the file or directory for which you want to enable auditing access.
4. Right-click the file or directory, and then choose **Properties**.

5. Choose the **Security** tab.
6. Click **Advanced**.
7. Choose the **Auditing** tab.
8. Perform the desired actions:

If you want to...	Do the following
Set up auditing for a new user or group	<ol style="list-style-type: none"> 1. Choose Add. 2. In the Enter the object name to select box, type the name of the user or group that you want to add. 3. Choose OK.
Remove auditing from a user or group	<ol style="list-style-type: none"> 1. In the Enter the object name to select box, select the user or group that you want to remove. 2. Choose Remove. 3. Choose OK. 4. Skip the rest of this procedure.
Change auditing for a user or group	<ol style="list-style-type: none"> 1. In the Enter the object name to select box, choose the user or group that you want to change. 2. Choose Edit. 3. Choose OK.

If you are setting up auditing on a user or group or changing auditing on an existing user or group, the **Auditing Entry for *object*** box opens.

9. In the **Apply to** box, select how you want to apply this auditing entry.

If you are setting up auditing on a single file, the **Apply to** box is not active, as it defaults to This object only.

10. In the **Access** box, select what you want audited and whether you want to audit successful events, failure events, or both.
 - To audit successful events, choose the **Success** box.
 - To audit failure events, choose the **Failure** box.

Choose the actions that you need to monitor to meet your security requirements. For more information about these auditable events, see your Windows documentation. You can audit the following events:

- Full control
 - Traverse folder / execute file
 - List folder / read data
 - Read attributes
 - Read extended attributes
 - Create files / write data
 - Create folders / append data
 - Write attributes
 - Write extended attributes
 - Delete subfolders and files
 - Delete
 - Read permissions
 - Change permissions
 - Take ownership
11. If you do not want the auditing setting to propagate to subsequent files and folders of the original container, choose the **Apply these auditing entries to objects and/or containers within this container only** box.
 12. Choose **Apply**.
 13. After you finish adding, removing, or editing auditing entries, choose **OK**.

The **Auditing Entry for *object*** box closes.

14. In the **Auditing** box, choose the inheritance settings for this folder. Choose only the minimal level that provides the auditing events that meet your security requirements.

You can choose one of the following:

- Choose the **Include inheritable auditing entries from this object's parent** box.
- Choose the **Replace all existing inheritable auditing entries on all descendants with inheritable auditing entries from this object** box.

- Choose both boxes.
- Choose neither box.

If you are setting SACLs on a single file, the **Replace all existing inheritable auditing entries on all descendants with inheritable auditing entries from this object** box is not present in the **Auditing** box.

15. Choose **OK**.

To configure NTFS audit policies (ONTAP CLI)

By using the ONTAP CLI, you can configure NTFS audit policies without needing to connect to the data using an SMB share on a Windows client.

- You can configure NTFS audit policies by using the [vserver security file-directory](#) command family.

For example, the following command applies a security policy named p1 to the SVM named vs0.

```
vserver security file-directory apply -vserver vs0 -policy-name p1
```

Audit policies on UNIX security-style files and directories

You configure auditing for UNIX security-style files and directories by adding audit ACEs (access control expressions) to NFS v4.x ACLs (access control lists). This allows you to monitor certain NFS file and directory access events for security purposes.

Note

For NFS v4.x, both discretionary and system ACEs are stored in the same ACL. Therefore, you must be careful when adding audit ACEs to an existing ACL to avoid overwriting and losing an existing ACL. The order in which you add the audit ACEs to an existing ACL does not matter.

To configure UNIX audit policies

1. Retrieve the existing ACL for the file or directory by using the `nfs4_getfacl` or equivalent command.
2. Append the desired audit ACEs.
3. Apply the updated ACL to the file or directory by using the `nfs4_setfacl` or equivalent command.

This example uses the `-a` option to give a user (named `testuser`) read permissions to the file named `file1`.

```
nfs4_setfacl -a "A::testuser@example.com:R" file1
```

Viewing audit event logs

You can view audit event logs saved in the EVTX or XML file formats.

- **EVTX file format** – You can open the converted EVTX audit event logs as saved files using Microsoft Event Viewer.

There are two options that you can use when viewing event logs using Event Viewer:

- **General view:** Information that is common to all events is displayed for the event record. The event-specific data for the event record is not displayed. You can use the detailed view to display event-specific data.
- **Detailed view:** A friendly view and a XML view are available. The friendly view and the XML view display both the information that is common to all events and the event-specific data for the event record.
- **XML file format** – You can view and process XML audit event logs on third-party applications that support the XML file format. XML viewing tools can be used to view the audit logs provided you have the XML schema and information about definitions for the XML fields.

Scaling SSD storage capacity and provisioned IOPS

When you need additional storage for the active portion of your data set, you can increase the solid state drive (SSD) storage capacity of your Amazon FSx for NetApp ONTAP file system. You

can do so by using the Amazon FSx console, Amazon FSx API, or Amazon Command Line Interface (Amazon CLI).

You can also change the provisioned SSD IOPS for your file system, either when you increase the primary SSD storage capacity or as an independent action. For more information on scaling a file system's primary SSD storage capacity and the amount of provisioned IOPS, see [Updating file system SSD storage and IOPS](#).

Managing throughput capacity

FSx for ONTAP configures throughput capacity when you create the file system. You can modify your scale-up file system's throughput capacity at any time, but you can't modify the throughput capacity of your scale-out file system. Keep in mind that your file system requires a specific configuration to achieve the maximum amount of throughput capacity. For example, to provision 4 GBps of throughput capacity for a scale-up file system, your file system requires a configuration with a minimum of 5,120 GiB of SSD storage capacity and 160,000 SSD IOPS. For more information, see [Impact of throughput capacity on performance](#).

Throughput capacity is one factor that determines the speed at which the file server that's hosting the file system can serve the file data. Higher levels of throughput capacity come with higher levels of network, disk read I/O operations per second (IOPS), and data caching capacity on the file server. For more information, see [Performance](#).

When you modify your file system's throughput capacity, Amazon FSx switches out the file server that's powering your file system. Both Single-AZ and Multi-AZ file systems experience an automatic failover and failback during this process, which typically takes a few minutes to complete. The failover and failback processes are transparent to NFS (Network File Sharing), SMB (Server Message Block), and iSCSI (Internet Small Computer Systems Interface) clients, allowing your workloads to continue running without interruption or manual intervention. You are billed for the new amount of throughput capacity once it's available to your file system.

Note

To ensure data integrity during maintenance activity, FSx for ONTAP closes all opportunistic locks and completes any pending write operations to the underlying storage volumes that are hosting your file system before maintenance begins. During a scheduled file system maintenance window, system modifications (such as modifications to your throughput capacity) may be delayed. System maintenance can cause these

changes to queue up until they are processed. For more information, see [the section called “Maintenance windows”](#).

Topics

- [When to modify throughput capacity](#)
- [How concurrent throughput and storage scaling requests are handled](#)
- [How to modify throughput capacity](#)
- [Monitoring throughput capacity changes](#)

When to modify throughput capacity

Amazon FSx integrates with Amazon CloudWatch, which helps you to monitor your file system's ongoing throughput usage levels. The throughput and IOPS performance that you can drive through your file system depends on your specific workload's characteristics, in addition to your file system's throughput capacity. As a rule, you should provision enough throughput capacity to support your workload's read throughput plus twice your workload's write throughput. You can use CloudWatch metrics to determine which of these dimensions to change to improve performance. For more information, see [the section called “How to use FSx for ONTAP CloudWatch metrics”](#).

Note

You can't modify the throughput capacity for scale-out file systems.

How concurrent throughput and storage scaling requests are handled

You can request a throughput capacity update just before an SSD storage capacity and provisioned IOPS update workflow begins or while it is in progress. The sequence of how Amazon FSx handles the two requests is as follows:

- If you submit an SSD/IOPS update and throughput capacity update at the same time, both requests are accepted. The SSD/IOPS update is prioritized before the throughput capacity update.
- If you submit a throughput capacity update while an SSD/IOPS update is in progress, the throughput capacity update request is accepted and queued to occur after the SSD/IOPS update.

The throughput capacity update starts after SSD/IOPS is updated (new values are available) and during the optimization step. This typically takes less than 10 minutes.

- If you submit a SSD/IOPS update while a throughput capacity update is in progress, the SSD/IOPS storage update request is accepted and queued to start after the throughput capacity update has completed (new throughput capacity is available). This typically takes 20 minutes.

For more information on SSD storage and provisioned IOPS updates, see [Managing storage capacity](#).

How to modify throughput capacity

You can modify a file system's throughput capacity using the Amazon FSx console, the Amazon Command Line Interface (Amazon CLI), or the Amazon FSx API.

To modify a file system's throughput capacity (console)

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. Navigate to **File systems**, and choose the ONTAP file system that you want to increase the throughput capacity for.
3. For **Actions**, choose **Update throughput capacity**. Or, in the **Summary** panel, choose **Update** next to the file system's **Throughput capacity**.
4. Choose the new value for **Throughput capacity** from the list.

Note

You can change the throughput capacity for any FSx for ONTAP file system. However, only file systems created on or after December 9, 2021 can support a throughput capacity of 128 MB/s or 256 MB/s.

5. Choose **Update** to initiate the throughput capacity update.
6. You can monitor the update progress on the **File systems** detail page, on the **Updates** tab.

You can monitor the progress of the update by using the Amazon FSx console, the Amazon CLI, and the API. For more information, see [Monitoring throughput capacity changes](#).

To modify a file system's throughput capacity (CLI)

To modify a file system's throughput capacity, use the Amazon CLI command [update-file-system](#). Set the following parameters:

- `--file-system-id` to the ID of the file system that you are updating.
- `ThroughputCapacity` to the desired value to update the file system to.

You can monitor the progress of the update by using the Amazon FSx console, the Amazon CLI, and the API. For more information, see [Monitoring throughput capacity changes](#).

Monitoring throughput capacity changes

You can monitor the progress of a throughput capacity modification using the Amazon FSx console, the API, and the Amazon CLI.

Monitoring throughput capacity changes in the console

On the **Updates** tab in the **File system details** window, you can view the 10 most recent update actions for each update action type.

For throughput capacity update actions, you can view the following information.

Update type

Supported types are **Throughput capacity**, **Storage capacity**, and **Storage optimization**.

Target value

The desired value to change the file system's throughput capacity to.

Status

The current status of the update. For throughput capacity updates, the possible values are as follows:

- **Pending** – Amazon FSx has received the update request, but has not started processing it.
- **In progress** – Amazon FSx is processing the update request.
- **Completed** – The throughput capacity update completed successfully.
- **Failed** – The throughput capacity update failed. Choose the question mark (?) to see details on why the throughput update failed.

Request time

The time when Amazon FSx received the update request.

Monitoring changes with the Amazon CLI and API

You can view and monitor file system throughput capacity modification requests using the [describe-file-systems](#) CLI command and the [DescribeFileSystems](#) API action. The `AdministrativeActions` array lists the 10 most recent update actions for each administrative action type. When you modify a file system's throughput capacity, a `FILE_SYSTEM_UPDATE` administrative action is generated.

The following example shows the response excerpt of a `describe-file-systems` CLI command. The file system has a throughput capacity of 128 MB/s, and a target throughput capacity of 256 MB/s.

```
.  
. .  
  "ThroughputCapacity": 128,  
  "AdministrativeActions": [  
    {  
      "AdministrativeActionType": "FILE_SYSTEM_UPDATE",  
      "RequestTime": 1581694764.757,  
      "Status": "PENDING",  
      "TargetFileSystemValues": {  
        "OntapConfiguration": {  
          "ThroughputCapacity": 256  
        }  
      }  
    }  
  ]
```

When Amazon FSx processes the action successfully, the status changes to `COMPLETED`. The new throughput capacity is then available to the file system, and shows in the `ThroughputCapacity` property. This is shown in the following response excerpt of a **describe-file-systems** CLI command.

```
.  
. .  
.
```

```
"ThroughputCapacity": 256,
"AdministrativeActions": [
  {
    "AdministrativeActionType": "FILE_SYSTEM_UPDATE",
    "RequestTime": 1581694764.757,
    "Status": "COMPLETED",
    "TargetFileSystemValues": {
      "OntapConfiguration": {
        "ThroughputCapacity": 256
      }
    }
  }
]
```

If the throughput capacity modification fails, the status changes to FAILED, and the `FailureDetails` property provides information about the failure.

Optimizing performance with Amazon FSx maintenance windows

As a fully-managed service, FSx for ONTAP regularly performs maintenance on and updates to your file system. This maintenance has no impact for most workloads. For workloads that are performance-sensitive, on rare occasions you may notice a brief (<60 seconds) impact on performance when maintenance is occurring; Amazon FSx enables you to use the maintenance window to control when any such potential maintenance activity occurs.

Patching occurs infrequently, typically once every several weeks. For scale-up file systems, patching typically only requires 30 minutes from the start of your maintenance window. For scale-out file systems, patching requires up to 90 minutes from the start of your maintenance window. During these few minutes, your file systems automatically fail over and fail back. You choose the maintenance window during file system creation. If you have no time preference, then a 30-minute start time is assigned.

FSx for ONTAP allows you to adjust your maintenance window as needed to accommodate your workload and operational requirements. You can move your maintenance window as frequently as required, provided that a maintenance window is scheduled at least once every 14 days. If a patch is released and you haven't scheduled a maintenance window within 14 days, FSx for ONTAP will proceed with maintenance on the file system to ensure its security and reliability.

Note

To ensure data integrity during maintenance activity, FSx for ONTAP closes all opportunistic locks and completes any pending write operations to the underlying storage volumes that are hosting your file system before maintenance begins.

You can use the Amazon FSx Management Console, Amazon CLI, Amazon API, or one of the Amazon SDKs to change the maintenance window for your file systems.

To change the weekly maintenance window (console)

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. Choose **File systems** in the left hand navigation column.
3. Choose the file system that you want to change the weekly maintenance window for. The **Summary** file system details page appears.
4. Choose **Administration** to display the file system administration **Settings** panel.
5. Choose **Update** to display the **Change maintenance window** window.
6. Enter the new day and time that you want the weekly maintenance window to start.
7. Choose **Save** to save your changes. The new maintenance start time is displayed in the file system administration **Settings** panel.

To change the weekly maintenance window using the [update-file-system](#) CLI command, see [To update a file system \(CLI\)](#).

Tag your Amazon FSx resources

To help you manage your file systems and other Amazon FSx resources, you can assign your own metadata to each resource in the form of *tags*. With tags, you can categorize your Amazon resources in different ways, for example, by purpose, owner, or environment. This categorization is useful when you have many resources of the same type—you can quickly identify a specific resource based on the tags that you've assigned to it. This topic describes tags and shows you how to create them.

Topics

- [Tag basics](#)

- [Tagging your resources](#)
- [Copying tags to backups](#)
- [Tag restrictions](#)
- [Permissions and tagging](#)

Tag basics

A *tag* is a label that you assign to an Amazon resource. Each tag consists of two parts that you define:

- A *tag key* (for example, `CostCenter`, `Environment`, or `Project`). Tag keys are case sensitive.
- A *tag value* (for example, `111122223333` or `Production`). Like tag keys, tag values are case sensitive. Tag values are optional.

You can use tags to categorize your Amazon resources in different ways, such as, by purpose, owner, or environment. For example, you could define a set of tags for your account's Amazon FSx file systems that helps you track each instance's owner and stack level.

We recommend that you devise a set of tag keys that meets your needs for each resource type. Using a consistent set of tag keys makes it easier for you to manage your resources. You can search and filter the resources based on the tags that you add. For more information about how to implement an effective resource tagging strategy, see [Tagging Amazon resources](#) in the *Amazon Web Services General Reference*.

Some tagging behaviors to keep in mind:

- Tags don't have any semantic meaning to Amazon FSx and are interpreted strictly as a string of characters.
- Tags are not automatically assigned to your resources.
- You can edit tag keys and values, and you can remove tags from a resource at any time.
- You can set the value of a tag to an empty string, but you can't set the value of a tag to `null`.
- If you add a tag that has the same key as an existing tag on that resource, the new value overwrites the old value.

- If you delete a resource, any tags for the resource are also deleted.
- If you're using the Amazon FSx API, the Amazon Command Line Interface (Amazon CLI), or an Amazon SDK, you can do the following:
 - You can use the `TagResource` API action to apply tags to existing resources.
 - For some resource-creating actions, you can specify tags for a resource when the resource is created. By tagging resources at the time of creation, you can eliminate the need to run custom tagging scripts after resource creation.

If tags cannot be applied during resource creation, Amazon FSx rolls back the resource creation process. This behavior helps ensure that resources are either created with tags or not created at all, and that no resources are left untagged at any time.

Note

Certain Amazon Identity and Access Management (IAM) permissions are required for users to tag resources on creation. For more information, see [Grant permission to tag resources during creation](#).

Tagging your resources

You can tag Amazon FSx resources that exist in your account. If you're using the Amazon FSx console, you can apply tags to resources by using the **Tags** tab on the relevant resource screen. When you create resources, you can apply the **Name** key with a value, and you can apply tags of your choice when creating a new file system. However, even though the console organizes resources according to the **Name** key, this key doesn't have any semantic meaning to the Amazon FSx service.

To implement granular control over the users and groups that can tag resources on creation, you can apply tag-based resource-level permissions in your IAM policies to the Amazon FSx API actions that support tagging on creation. By using such permissions in your policies, you get the following benefits:

- Your resources are properly secured from creation.
- Because tags are applied immediately to your resources, any tag-based resource-level permissions controlling the use of resources are immediately effective.
- Your resources can be tracked and reported on more accurately.

- You can enforce the use of tagging on new resources, and control which tag keys and values are set on your resources.

To control which tag keys and values are set on your existing resources, you can apply resource-level permissions to the `TagResource` and `UntagResource` Amazon FSx API actions in your IAM policies.

For more information about the permissions required to tag Amazon FSx resources at creation, see [Grant permission to tag resources during creation](#).

For more information about using tags to restrict access to Amazon FSx resources in IAM policies, see [Using tags to control access to your Amazon FSx resources](#).

For information about tagging your resources for billing, see [Using cost allocation tags](#) in the *Amazon Billing User Guide*.

Copying tags to backups

When you create or update a volume in the Amazon FSx API or Amazon CLI, you can enable `CopyTagsToBackups` to automatically copy any tags from your volumes to backups.

Note

If you specify tags while creating a user-initiated backup (including the name tag when you create a backup using the Amazon FSx console), tags are *not* copied from the volume even if you've enabled `CopyTagsToBackups`.

For more information about backups, see [Working with backups](#). For more information about enabling `CopyTagsToBackups`, see [To create a volume \(CLI\)](#) and [To update a volume's configuration \(CLI\)](#) in the *Amazon FSx for NetApp ONTAP User Guide* or [CreateVolume](#) and [UpdateVolume](#) in the *Amazon FSx for NetApp ONTAP API Reference*.

Tag restrictions

The following basic restrictions apply to tags:

- The maximum number of tags per resource is 50.

- The maximum key length is 128 Unicode characters in UTF-8.
- The maximum value length is 256 Unicode characters in UTF-8.
- The allowed characters are letters, numbers, and spaces representable in UTF-8, and the following characters: + - (hyphen) = . _ (underscore) : / @.
- For each resource, each tag key must be unique, and each tag key can have only one value.
- Tag keys and values are case sensitive.
- The `aws :` prefix is reserved for Amazon use. If a tag has a tag key with this prefix, you can't edit or delete the tag's key or value. Tags with the `aws :` prefix do not count against your tags per resource limit.

You can't delete a resource based solely on its tags; you must specify the resource identifier. For example, to delete a file system that you tagged with a tag key called `DeleteMe`, you must use the `DeleteFileSystem` action with the file system resource identifier, such as `fs-1234567890abcdef0`.

When you tag public or shared resources, the tags that you assign are available only to your Amazon Web Services account; no other Amazon Web Services account has access to those tags. For tag-based access control to shared resources, each Amazon Web Services account must assign its own set of tags to control access to the resource.

Permissions and tagging

For more information about the permissions required to tag Amazon FSx resources at creation, see [Grant permission to tag resources during creation](#).

For more information about using tags to restrict access to Amazon FSx resources in IAM policies, see [Using tags to control access to your Amazon FSx resources](#).

Managing FSx for ONTAP resources using NetApp applications

In addition to the Amazon Web Services Management Console, Amazon CLI, and Amazon API and SDKs, you can also use these NetApp management tools and applications to manage your FSx for ONTAP resources:

Topics

- [Signing up for a NetApp account](#)

- [Using NetApp BlueXP](#)
- [Using the NetApp ONTAP CLI](#)
- [Using the ONTAP REST API](#)

Important

Amazon FSx periodically syncs with ONTAP to ensure consistency. If you create or modify volumes using NetApp applications, it may take up to several minutes for these changes to be reflected in the Amazon Web Services Management Console, Amazon CLI, API and SDKs.

Signing up for a NetApp account

In order to download some NetApp software, such as BlueXP, SnapCenter, and the ONTAP Antivirus connector, you need to have a NetApp account. To sign up for a NetApp account, perform the following steps:

1. Go to the [NetApp User Registration](#) page and register for a new NetApp user account.
2. Complete the form(s) with your information. Be sure to select the **NetApp Customer/End User** access level. In the **SERIAL NUMBER** field, copy and paste the File System ID for your FSx for ONTAP file system. See the following example:

USER ACCESS LEVEL

- Guest User NetApp Customer / End User
 NetApp Reseller / Service Provider / System Integrator / Partner

Product Information (Optional)

Please enter a Serial Number or System ID to help us validate your access level.

Please note: Not providing a Serial Number or System ID may delay processing of your request.

SERIAL NUMBER

(Either a NetApp hardware Serial Number, often located on back of unit; or a NetApp software Serial Number.)

OR

SYSTEM ID

(Run a "sysconfig -a" command on your NetApp product. The output should list the System ID.)

NETAPP TOKEN

What to expect after you register

Customers with existing NetApp products will have their NSS account leveled-up to **Customer Level** access within one business day. Customers new to NetApp will be onboarded using standard business practices, in addition to having their NSS account leveled-up to Customer Level access. Providing the File System ID helps expedite this process. You can check the status of your NSS account by logging into mysupport.netapp.com and navigating to the **Welcome** page. The access level of your account should be **Customer Access**.

Using NetApp BlueXP

NetApp BlueXP is a unified control plane that simplifies management experiences for storage and data services across on-premises and cloud environments. BlueXP provides a centralized user interface to manage, monitor, and automate ONTAP deployments in Amazon and on premises. For

more information, see the [NetApp BlueXP documentation](#) and the [NetApp BlueXP for Amazon FSx for NetApp ONTAP](#) documentation.

Note

NetApp BlueXP isn't supported for scale-out file systems.

Using NetApp System Manager with BlueXP

You can manage your Amazon FSx for NetApp ONTAP file systems using System Manager directly from BlueXP. BlueXP lets you use the same System Manager interface that you're accustomed to using, so you can manage your hybrid multi-cloud infrastructure from a single control plane. You also have access to BlueXP's other functionality. For more information, see the [System Manager integration with BlueXP](#) topic in the NetApp ONTAP documentation.

Note

NetApp System Manager isn't supported for scale-out file systems.

Using the NetApp ONTAP CLI

You can manage your Amazon FSx for NetApp ONTAP resources using the NetApp ONTAP CLI. You can manage resources at the file system (analogous to NetApp ONTAP cluster) level, and at the SVM level.

Managing file systems with the ONTAP CLI

You can run ONTAP CLI commands on your FSx for ONTAP file system, analogous to running them on a NetApp ONTAP cluster. You access the ONTAP CLI on your file system by establishing a secure shell (SSH) connection to the file system's management endpoint, logging in with the `fsxadmin` username and password. You have the option to set the password when you create file system using the custom create flow or using the Amazon CLI. If you created the file system using the Quick create option, the `fsxadmin` password was not set, so you'll to set one in order to log in to the ONTAP CLI. For more information, see [Updating a file system](#). You can find the **DNS name** and **IP address** of your file system's management endpoint in the Amazon FSx console, in the **Administration** tab of the FSx for ONTAP file system details page, shown in the following graphic.

The screenshot shows the 'Administration' tab of the ONTAP console. Under 'ONTAP administration', there are several configuration items:

- Management endpoint - DNS name:** `management.fs-08fc3405e03933af0.fsx.us-east-2.aws.com` (indicated by a blue arrow)
- Management endpoint - IP address:** `198.19.255.184` (indicated by a blue arrow)
- Inter-cluster endpoint - DNS name:** `intercluster.fs-08fc3405e03933af0.fsx.us-east-2.aws.com`
- Inter-cluster endpoint - IP address:** `172.31.32.114` and `172.31.2.110`
- Service account username:** `fsxadmin`
- Service account password:** `<INTENTIONALLY REDACTED>`
- Update button:** A button labeled 'Update' is located to the right of the password field.

To connect to the file system's management endpoint with SSH, use the `fsxadmin` user and password. You can SSH into the file system's management endpoint IP address or DNS name from a client that is in the same VPC as the file system, as in the following examples.

```
ssh fsxadmin@file-system-management-endpoint-ip-address
```

The SSH command with sample values:

```
ssh fsxadmin@198.51.100.0
```

The SSH command using the management endpoint DNS name:

```
ssh fsxadmin@file-system-management-endpoint-dns-name
```

The SSH command using a sample DNS name:

```
$ ssh fsxadmin@management.fs-0abcdef123456789.fsx.us-east-2.aws.com
Password: fsxadmin-password
```

This is your first recorded login.
FsxId0abcdef123456789::>

Scope of ONTAP CLI commands available to fsxadmin

The fsxadmin's administrative view is at the file system level, which includes all SVMs and volumes in the file system. The fsxadmin role performs the role of the ONTAP cluster administrator. Because Amazon FSx for NetApp ONTAP file systems are fully managed, the fsxadmin role can run a subset of the available ONTAP CLI commands.

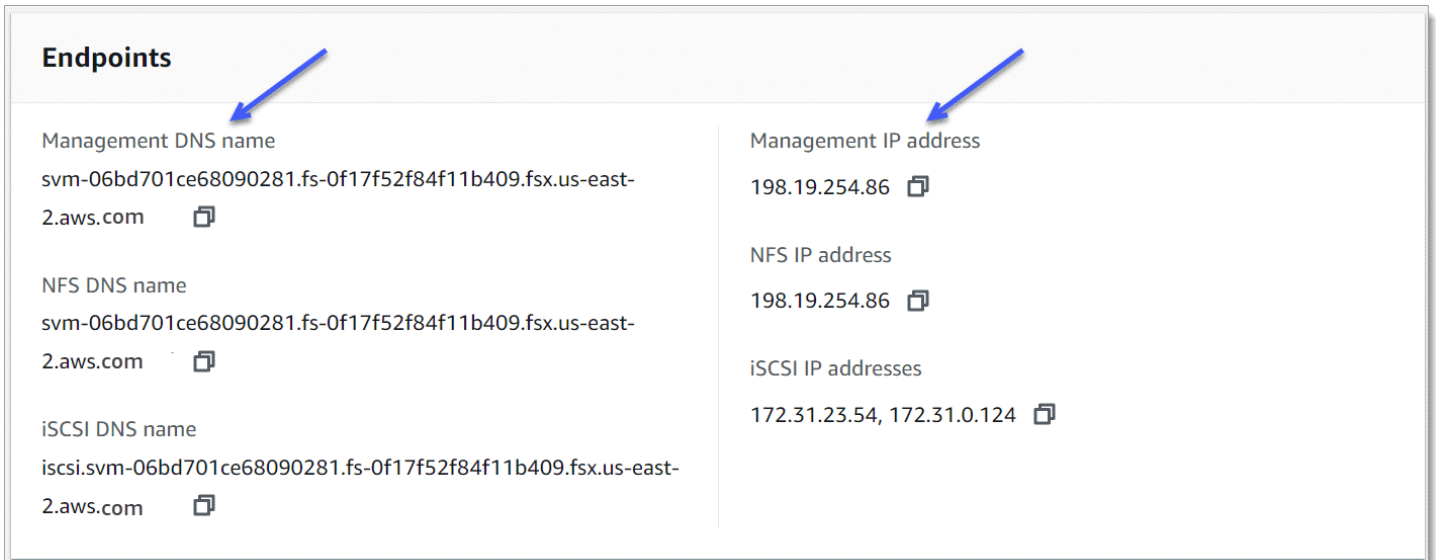
To see a list of the commands that fsxadmin can run, use the following [security login role show](#) ONTAP CLI command:

```
FsxId0abc123def456::> security login role show -role fsxadmin -access !none
      Role          Command/          Access
Vserver  Name          Directory          Query Level
-----
FsxId0abcdef123456789
      fsxadmin    application          all
      cluster application-record          all
      cluster date show          readonly
      cluster ha modify          readonly
      cluster ha show          readonly
      cluster identity modify          readonly
      cluster identity show          readonly
      cluster log-forwarding          -port !55555 all
      cluster modify          readonly
      cluster peer          all
      cluster show          readonly
      cluster statistics show          readonly
      cluster time-service ntp server create          readonly
      cluster time-service ntp server delete          readonly
      cluster time-service ntp server modify          readonly
      cluster time-service ntp server show          readonly
      debug network tcpdump          -ipspace !Cluster all
      debug san lun          all
      df          -vserver !FsxId* -vserver !Cluster readonly
      echo          all
      event catalog show          readonly
      event config          all
.
.
.
```


363 entries were displayed.

Managing SVMs with the ONTAP CLI

You can access the ONTAP CLI on your SVM by establishing a secure shell (SSH) connection to the SVM's management endpoint using either the `fsxadmin` or the `vsadmin` user name and password. You can find the SVM's management endpoint **DNS name** and **IP address** in the Amazon FSx console, in the **Endpoints** panel of the **Storage virtual machines** details page, shown in the following graphic.



To connect to the SVM's management endpoint with SSH, you can use either the `vsadmin` or `fsxadmin` username and password. If you did not set a password for the `vsadmin` user when the SVM was created, you can set the `vsadmin` password at anytime. For more information, see [Updating a storage virtual machine](#). You can SSH into the SVM from a client that is in the same VPC as the file system, using the management endpoint IP address or DNS name.

```
ssh vsadmin@svm-management-endpoint-ip-address
```

The command with sample values:

```
ssh vsadmin@198.51.100.10
```

The SSH command using the management endpoint DNS name:

```
ssh vsadmin@svm-management-endpoint-dns-name
```

The SSH command using a sample DNS name:

```
ssh vsadmin@management.svm-abcdef01234567892fs-0abcdef123456789.fsx.us-east-2.aws.com
```

Password: *vsadmin-password*

This is your first recorded login.

```
FsxId0abcdef123456789::>
```

Amazon FSx for NetApp ONTAP supports the NetApp ONTAP CLI commands.

For a complete reference of NetApp ONTAP CLI commands, see the [ONTAP Commands: Manual Page Reference](#).

Using the ONTAP REST API

When accessing your FSx for ONTAP file system using the ONTAP REST API using the `fsxadmin` credentials, do one of the following:

- Disable TLS validation.

Or

- Trust the Amazon certificate authorities (CAs) – The certificate bundle for the CAs in each region can be found at the follow URLs:
 - <https://fsx-aws-certificates.s3.amazonaws.com/bundle-aws-region.pem> for Public Amazon Web Services Regions
 - <https://fsx-aws-us-gov-certificates.s3.us-gov-west-1.amazonaws.com/bundle-aws-region.pem> for AmazonGovCloud Regions
 - <https://fsx-aws-cn-certificates.s3.cn-north-1.amazonaws.com.cn/bundle-aws-region.pem> for Amazon China Regions

For a complete reference of NetApp ONTAP REST API commands, see the [NetApp ONTAP REST API Online Reference](#).

Security in Amazon FSx for NetApp ONTAP

Cloud security at Amazon is the highest priority. As an Amazon customer, you benefit from data centers and network architectures that are built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between Amazon and you. The [shared responsibility model](#) describes this as security *of* the cloud and security *in* the cloud:

- **Security of the cloud** – Amazon is responsible for protecting the infrastructure that runs Amazon services in the Amazon Web Services Cloud. Amazon also provides you with services that you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the [Amazon Compliance Programs](#). To learn about the compliance programs that apply to Amazon FSx for NetApp ONTAP, see [Amazon Services in Scope by Compliance Program](#).
- **Security in the cloud** – Your responsibility is determined by the Amazon service that you use. You are also responsible for other factors including the sensitivity of your data, your company's requirements, and applicable laws and regulations.

This documentation helps you understand how to apply the shared responsibility model when using Amazon FSx. The following topics show you how to configure Amazon FSx to meet your security and compliance objectives. You also learn how to use other Amazon services that help you to monitor and secure your Amazon FSx resources.

Topics

- [Data protection in Amazon FSx for NetApp ONTAP](#)
- [Identity and access management for Amazon FSx for NetApp ONTAP](#)
- [Amazon managed policies for Amazon FSx](#)
- [File System Access Control with Amazon VPC](#)
- [Compliance Validation for Amazon FSx for NetApp ONTAP](#)
- [Amazon FSx for NetApp ONTAP and interface VPC endpoints \(Amazon PrivateLink\)](#)
- [Resilience in Amazon FSx for NetApp ONTAP](#)
- [Infrastructure security in Amazon FSx for NetApp ONTAP](#)
- [Use NetApp ONTAP Vscan with FSx for ONTAP](#)

- [Roles and users in Amazon FSx for NetApp ONTAP](#)

Data protection in Amazon FSx for NetApp ONTAP

The Amazon [shared responsibility model](#) applies to data protection in Amazon FSx for NetApp ONTAP. As described in this model, Amazon is responsible for protecting the global infrastructure that runs all of the Amazon Web Services Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. You are also responsible for the security configuration and management tasks for the Amazon Web Services that you use. For more information about data privacy, see the [Data Privacy FAQ](#).

For data protection purposes, we recommend that you protect Amazon Web Services account credentials and set up individual users with Amazon IAM Identity Center or Amazon Identity and Access Management (IAM). That way, each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.
- Use SSL/TLS to communicate with Amazon resources. We require TLS 1.2 and recommend TLS 1.3.
- Set up API and user activity logging with Amazon CloudTrail.
- Use Amazon encryption solutions, along with all default security controls within Amazon Web Services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing sensitive data that is stored in Amazon S3.
- If you require FIPS 140-2 validated cryptographic modules when accessing Amazon through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see [Federal Information Processing Standard \(FIPS\) 140-2](#).

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form text fields such as a **Name** field. This includes when you work with Amazon FSx or other Amazon Web Services using the console, API, Amazon CLI, or Amazon SDKs. Any data that you enter into tags or free-form text fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

Data encryption in FSx for ONTAP

Amazon FSx for NetApp ONTAP supports encryption of data at rest and encryption of data in transit. Encryption of data at rest is automatically enabled when creating an Amazon FSx file system. Amazon FSx for NetApp ONTAP supports Kerberos-based encryption in transit over the NFS and SMB protocols if you're accessing data in a Storage Virtual Machine (SVM) that's joined to an Active Directory or to a domain using the Lightweight Directory Access Protocol (LDAP).

When to use encryption

If your organization is subject to corporate or regulatory policies that require encryption of data and metadata at rest, your data is automatically encrypted at rest. We also recommend that you enable encryption of data in transit by mounting your file system using encryption of data in transit.

For more information about data encryption with Amazon FSx for NetApp ONTAP, see [Encryption of data at rest](#) and [Encrypting data in transit](#).

Encryption of data at rest

All Amazon FSx for NetApp ONTAP file systems are encrypted at rest with keys managed using Amazon Key Management Service (Amazon KMS). Data is automatically encrypted before being written to the file system, and automatically decrypted as it is read. These processes are handled transparently by Amazon FSx, so you don't have to modify your applications.

Amazon FSx uses an industry-standard AES-256 encryption algorithm to encrypt Amazon FSx data and metadata at rest. For more information, see [Cryptography Basics](#) in the *Amazon Key Management Service Developer Guide*.

Note

The Amazon key management infrastructure uses Federal Information Processing Standards (FIPS) 140-2 approved cryptographic algorithms. The infrastructure is consistent with National Institute of Standards and Technology (NIST) 800-57 recommendations.

How Amazon FSx uses Amazon KMS

Amazon FSx integrates with Amazon KMS for key management. Amazon FSx uses KMS keys to encrypt your file system. You choose the KMS key used to encrypt and decrypt file systems (both

data and metadata). You can enable, disable, or revoke grants on this KMS key. This KMS key can be one of the two following types:

- **Amazon-managed KMS key** – This is the default KMS key, and it's free to use.
- **Customer-managed KMS key** – This is the most flexible KMS key to use, because you can configure its key policies and grants for multiple users or services. For more information on creating KMS keys, see [Creating Keys](#) in the *Amazon Key Management Service Developer Guide*.

Important

Amazon FSx accepts only symmetric encryption KMS keys. You can't use asymmetric KMS keys with Amazon FSx.

If you use a customer-managed KMS key as your KMS key for file data encryption and decryption, you can enable key rotation. When you enable key rotation, Amazon KMS automatically rotates your key once per year. Additionally, with a customer-managed KMS key, you can choose when to disable, re-enable, delete, or revoke access to your KMS key at any time. For more information, see [Rotating Amazon KMS keys](#) and [Enabling and disabling keys](#) in the *Amazon Key Management Service Developer Guide*.

Amazon FSx key policies for Amazon KMS

Key policies are the primary way to control access to KMS keys. For more information on key policies, see [Using key policies in Amazon KMS](#) in the *Amazon Key Management Service Developer Guide*. The following list describes all the Amazon KMS-related permissions supported by Amazon FSx for encrypted at rest file systems:

- **kms:Encrypt** – (Optional) Encrypts plaintext into ciphertext. This permission is included in the default key policy.
- **kms:Decrypt** – (Required) Decrypts ciphertext. Ciphertext is plain text that has been previously encrypted. This permission is included in the default key policy.
- **kms:ReEncrypt** – (Optional) Encrypts data on the server side with a new Amazon KMS key, without exposing the plaintext of the data on the client side. The data is first decrypted and then re-encrypted. This permission is included in the default key policy.

- **kms:GenerateDataKeyWithoutPlaintext** – (Required) Returns a data encryption key encrypted under a KMS key. This permission is included in the default key policy under **kms:GenerateDataKey***.
- **kms:CreateGrant** – (Required) Adds a grant to a key to specify who can use the key and under what conditions. Grants are alternate permission mechanisms to key policies. For more information on grants, see [Using Grants](#) in the *Amazon Key Management Service Developer Guide*. This permission is included in the default key policy.
- **kms:DescribeKey** – (Required) Provides detailed information about the specified KMS key. This permission is included in the default key policy.
- **kms:ListAliases** – (Optional) Lists all of the key aliases in the account. When you use the console to create an encrypted file system, this permission populates the list of KMS keys. We recommend using this permission to provide the best user experience. This permission is included in the default key policy.

Encrypting data in transit

This topic explains the different options available for encrypting your file data while it is in-transit between an FSx for ONTAP file system and connected clients. It also provides guidance to help you choose which encryption method is best suited for your workflow.

All data flowing across Amazon Web Services Regions over the Amazon global network is automatically encrypted at the physical layer before it leaves Amazon secured facilities. All traffic between Availability Zones is encrypted. Additional layers of encryption, including those listed in this section, provide additional protections. For more information about how Amazon provides protection for data flowing across Amazon Web Services Regions, Available Zones, and instances, see [Encryption in transit](#) in the Amazon Elastic Compute Cloud User Guide for Linux Instances.

Amazon FSx for NetApp ONTAP supports the following methods for encrypting data in transit between FSx for ONTAP file systems and connected clients:

- Automatic Nitro-based encryption over all supported protocols and clients running on supported Amazon EC2 [Linux](#) and [Windows](#) instance types.
- Kerberos-based encryption over NFS and SMB protocols.
- IPsec-based encryption over NFS, iSCSI, and SMB protocols

All of the supported methods for encrypting data in transit use industry standard AES-256 cryptographic algorithms that provide enterprise strength encryption.

Topics

- [Choosing a method for encrypting data in transit](#)
- [Encrypting data in transit with Amazon Nitro System](#)
- [Encrypting data in-transit with Kerberos-based encryption](#)
- [Encrypting data in transit with IPsec encryption](#)
- [Enable SMB encryption of data in transit](#)
- [Configuring IPsec using PSK authentication](#)
- [Configuring IPsec using certificate authentication](#)

Choosing a method for encrypting data in transit

This section provides information that can help you decide which of the supported encryption in transit methods is best for your workflow. Refer back to this section as you explore the supported options described in detail in the sections that follow.

There are several factors to consider when choosing how you are going to encrypt data in transit between your FSx for ONTAP file system and connected clients. These factors include:

- The Amazon Web Services Region that your FSx for ONTAP file system is running in.
- The instance type that the client is running on.
- The location of the client accessing your file system.
- Network performance requirements.
- The data protocol you want to encrypt.
- If you are using Microsoft Active Directory.

Amazon Web Services Region

The Amazon Web Services Region that your file system is running in determines whether or not you can use Amazon Nitro-based encryption. Nitro-based encryption is available in the following Amazon Web Services Regions:

- US East (N. Virginia)
- US East (Ohio)

- US West (Oregon)
- Europe (Ireland)

Additionally, Nitro-based encryption is available for scale-out file systems in the Asia Pacific (Sydney) Amazon Web Services Region.

Client instance type

You can use Amazon Nitro-based encryption if the client accessing your file system is running on any of the supported Amazon EC2 Mac, [Linux](#) or [Windows](#) instance types, and your workflow meets all other requirements for using [Nitro-based encryption](#). There aren't any client instance type requirements for using Kerberos or IPsec encryption.

Client location

The location of the client accessing data with respect to the location of your file system impacts which in-transit encryption methods are available to use. You can use any of the supported encryption methods if the client and file system are located in the same VPC. The same is true if the client and file system are located in peered VPCs, as long as the traffic does not pass through a virtual network device or service, such as a transit gateway. Nitro-based encryption is not an available option if the client is not in the same or peered VPC, or if the traffic passes through a virtual network device or service.

Network performance

Using Amazon Nitro-based encrypted has no impact on network performance. This is because the supported Amazon EC2 instances utilize the offload capabilities of the underlying Nitro System hardware to automatically encrypt in-transit traffic between instances.

Using Kerberos or IPsec encryption has an impact on network performance. This is because both of these encryption methods are software-based, which requires the client and server to use compute resources to encrypt and decrypt in-transit traffic.

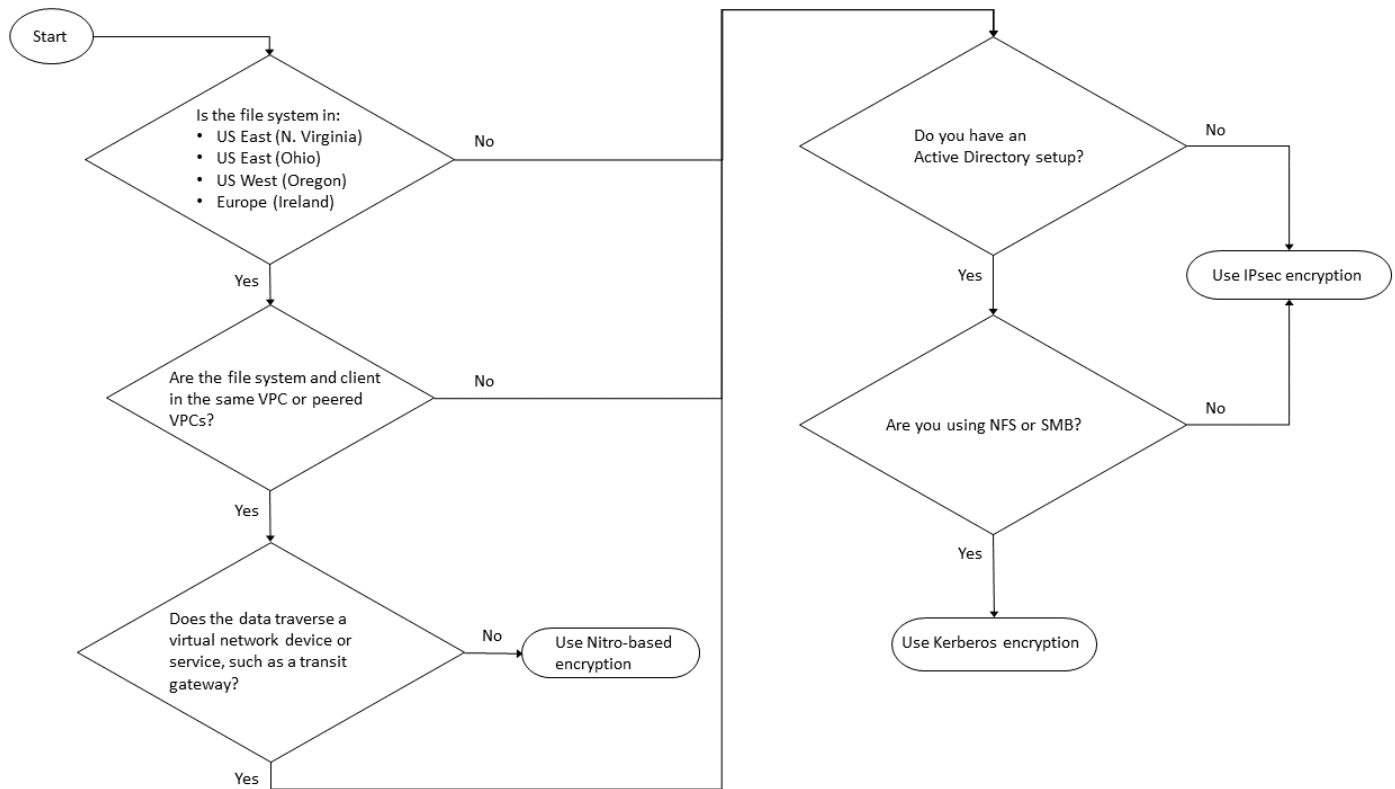
Data protocol

You can use Amazon Nitro-based encryption and IPsec encryption with all of the supported protocols – NFS, SMB, and iSCSI. You can use Kerberos encryption with the NFS and SMB protocols (with an Active Directory).

Active Directory

If you are using Microsoft Active Directory, you can use [Kerberos encryption](#) over the NFS and SMB protocols.

Use the following diagram to help you decide which encryption in-transit method to use.



IPsec encryption is the only option available when all of the following conditions apply to your workflow:

- You are using the NFS, SMB, or iSCSI protocol.
- Your workflow does not support using Amazon Nitro-based encryption.
- You are not using a Microsoft Active Directory domain.

Encrypting data in transit with Amazon Nitro System

With Nitro-based encryption, data in-transit is encrypted automatically when clients accessing your file systems are running on supported Amazon EC2 [Linux](#) or [Windows](#) instance types.

Using Amazon Nitro-based encryption has no impact on network performance. This is because the supported Amazon EC2 instances utilize the offload capabilities of the underlying Nitro System hardware to automatically encrypt in-transit traffic between instances.

Nitro-based encryption is enabled automatically when the supported client instance types are located in the same Amazon Web Services Region and in the same VPC or in a VPC peered with the

file system's VPC. Additionally, if the client is in a peered VPC, then data cannot traverse a virtual network device or service (such as a transit gateway) in order for Nitro-based encryption to be automatically enabled. For more information about Nitro-based encryption, see the Encryption in transit section of the Amazon EC2 User Guide for [Linux](#) or [Windows](#) instance types.

Nitro-based in-transit encryption is available for file systems created after November 28, 2022 in the following Amazon Web Services Regions:

- US East (N. Virginia)
- US East (Ohio)
- US West (Oregon)
- Europe (Ireland)

Additionally, Nitro-based encryption is available for scale-out file systems in the Asia Pacific (Sydney) Amazon Web Services Region.

For more information about the Amazon Web Services Regions where FSx for ONTAP is available, see [Amazon FSx for NetApp ONTAP Pricing](#).

For more information about the performance specifications for FSx for ONTAP file systems, see [Impact of throughput capacity on performance](#).

Encrypting data in-transit with Kerberos-based encryption

If you are using Microsoft Active Directory, you can use Kerberos-based encryption over the NFS and SMB protocols to encrypt data in-transit for child volumes of [SVMs that are joined to a Microsoft Active Directory](#).

Encrypting data in transit over NFS using Kerberos

Encryption of data in-transit using Kerberos is supported for NFSv3 and NFSv4 protocols. To enable encryption in transit using Kerberos for the NFS protocol, see [Using Kerberos with NFS for strong security](#) in the NetApp ONTAP Documentation Center.

Encrypting data in transit over SMB using Kerberos

Encrypting data in transit over the SMB protocol is supported on file shares that are mapped on a compute instance that supports SMB protocol 3.0 or newer. This includes all Microsoft Windows versions from Microsoft Windows Server 2012 and later, and Microsoft Windows 8 and later. When

enabled, FSx for ONTAP automatically encrypts data in transit using SMB encryption as you access your file system without the need for you to modify your applications.

FSx for ONTAP SMB supports 128 and 256 bit encryption, which is determined by the client session request. For descriptions of the different encryption levels, see the *Set the SMB server minimum authentication security level* section of [Manage SMB with the CLI](#) in the NetApp ONTAP Documentation Center.

Note

The client determines the encryption algorithm. Both NTLM and Kerberos authentication work with both 128 and 256 bit encryption. The FSx for ONTAP SMB Server accepts all standard Windows client requests, and the granular controls are handled by Microsoft Group Policy or Registry settings.

You use the ONTAP CLI to manage the encryption in transit settings on FSx for ONTAP SVMs and volumes. To access the NetApp ONTAP CLI, establish an SSH session on the SVM on which you are making encryption in transit settings, as described in [Managing SVMs with the ONTAP CLI](#).

For instructions on how to enable SMB encryption on an SVM or volume, see [Enable SMB encryption of data in transit](#).

Encrypting data in transit with IPsec encryption

FSx for ONTAP supports using the IPsec protocol in transport mode to ensure data is continuously secure and encrypted, while in-transit. IPsec offers end-to-end encryption of data in-transit between clients and FSx for ONTAP file systems for all supported IP traffic – NFS, iSCSI, and SMB protocols. With IPsec encryption, you establish an IPsec tunnel between an FSx for ONTAP SVM configured with IPsec enabled, and an IPsec client running on the connected client accessing the data.

We recommend that you use IPsec to encrypt data in transit over NFS, SMB, and iSCSI protocols when accessing your data from clients that do not support [Nitro-based encryption](#), and if your client and SVMs are not joined to an Active Directory, which is required for Kerberos-based encryption. IPsec encryption is the only option available for encrypting data in-transit for iSCSI traffic when your iSCSI client doesn't support Nitro-based encryption.

For IPsec authentication, you can use either pre-shared keys (PSKs) or certificates. If you are using a PSK, the IPsec client you use must support Internet Key Exchange version 2 (IKEv2) with a PSK.

The high level steps for configuring IPsec encryption on both FSx for ONTAP and the client are as follows:

1. Enable and configure IPsec on your file system.
2. Install and configure IPsec on your client
3. Configure IPsec for multiple client access

For more information about how to configure IPsec using PSK, see [Configure IP security \(IPsec\) over wire encryption](#) in the NetApp ONTAP documentation center.

For more information about how to configure IPsec using certificates, see [Configuring IPsec using certificate authentication](#).

Enable SMB encryption of data in transit

By default, when you create an SVM, SMB encryption is turned off. You can either enable SMB encryption required on individual shares, or on an SVM, which turns it on for all shares on that SVM.

Note

When SMB encryption required is enabled on an SVM or share, SMB clients that do not support encryption cannot connect to that SVM or share.

To require SMB encryption for incoming SMB traffic on an SVM

Use the following procedure to require SMB encryption on a SVM using the NetApp ONTAP CLI.

1. To connect to the SVM management endpoint with SSH, use user name `vsadmin` and the `vsadmin` password that you set when you created the SVM. If you did not set a `vsadmin` password, use user name `fsxadmin` and the `fsxadmin` password. You can SSH into the SVM from a client that is in the same VPC as the file system, using the management endpoint IP address or DNS name.

```
ssh vsadmin@svm-management-endpoint-ip-address
```

The command with sample values:

```
ssh vsadmin@198.51.100.10
```

The SSH command using the management endpoint DNS name:

```
ssh vsadmin@svm-management-endpoint-dns-name
```

The SSH command using a sample DNS name:

```
ssh vsadmin@management.svm-abcdef01234567892fs-08fc3405e03933af0.fsx.us-east-2.aws.com
```

Password: **vsadmin-password**

```
This is your first recorded login.
FsxIdabcdef01234567892::>
```

2. Use the [vserver cifs security modify](#) NetApp ONTAP CLI command to require SMB encryption for incoming SMB traffic to the SVM.

```
vserver cifs security modify -vserver vservice_name -is-smb-encryption-required true
```

3. To stop requiring SMB encryption for incoming SMB traffic, use the following command.

```
vserver cifs security modify -vserver vservice_name -is-smb-encryption-required false
```

4. To see the current is-smb-encryption-required setting on an SVM, use the [vserver cifs security show](#) NetApp ONTAP CLI command:

```
vserver cifs security show -vserver vs1 -fields is-smb-encryption-required

vserver  is-smb-encryption-required
-----  -
vs1      true
```

For more information about managing SMB encryption on an SVM, see [Configuring required SMB encryption on SMB servers for data transfers over SMB](#) in the NetApp ONTAP Documentation Center.

To enable SMB encryption on a volume

Use the following procedure to enable SMB encryption on a share using the NetApp ONTAP CLI.

1. Establish a secure shell (SSH) connection to the SVM's management endpoint as described in [Managing SVMs with the ONTAP CLI](#).
2. Use the following NetApp ONTAP CLI command to create a new SMB share and require SMB encryption when accessing this share.

```
vserver cifs share create -vserver vserver_name -share-name share_name -  
path share_path -share-properties encrypt-data
```

For more information, see [vserver cifs share create](#) in the NetApp ONTAP CLI Command man pages.

3. To require SMB encryption on an existing SMB share, use the following command.

```
vserver cifs share properties add -vserver vserver_name -share-name share_name -  
share-properties encrypt-data
```

For more information, see [vserver cifs share create](#) in the NetApp ONTAP CLI Command man pages.

4. To turn off SMB encryption on an existing SMB share, use the following command.

```
vserver cifs share properties remove -vserver vserver_name -share-name share_name -  
share-properties encrypt-data
```

For more information, see [vserver cifs share properties remove](#) in the NetApp ONTAP CLI Command man pages.

5. To see the current `is-smb-encryption-required` setting on an SMB share, use the following NetApp ONTAP CLI command:

```
vserver cifs share properties show -vserver vserver_name -share-name share_name -  
fields share-properties
```

If one of the properties returned by the command is the `encrypt-data` property, then that property specifies that SMB encryption must be used when accessing this share.

For more information, see [vserver cifs share properties show](#) in the NetApp ONTAP CLI Command man pages.

Configuring IPsec using PSK authentication

If you are using PSK for authentication, the steps for configuring IPsec encryption on both FSx for ONTAP and the client are as follows:

1. Enable and configure IPsec on your file system.
2. Install and configure IPsec on your client
3. Configure IPsec for multiple client access

For details on configuring IPsec using PSK, see [Configure IP security \(IPsec\) over wire encryption](#) in the NetApp ONTAP documentation center.

Configuring IPsec using certificate authentication

The following topics provide instructions for configuring IPsec encryption using certificate authentication on an FSx for ONTAP file system and a client running Libreswan IPsec. This solution uses Amazon Certificate Manager and Amazon Private Certificate Authority to create a private certificate authority and for generating the certificates.

The high-level steps for configuring IPsec encryption using certificate authentication on FSx for ONTAP file systems and connected clients are as follows:

1. Have a certificate authority in place for issuing certificates.
2. Generate and export CA certificates for the file system and client.
3. Install certificate and configure IPsec on the client instance.
4. Install certificate and configure IPsec on your file system.
5. Define the security policy database (SPD).
6. Configure IPsec for multiple client access.

Creating and installing CA certificates

For certificate authentication, you need to generate and install certificates from a certificate authority on your FSx for ONTAP file system and the clients that will access the data on your

file system. The following example uses Amazon Private Certificate Authority to set up a private certificate authority, and generate the certificates to install on the file system and the client. Using Amazon Private Certificate Authority, you can create an entirely Amazon hosted hierarchy of root and subordinate certificate authorities (CAs) for internal use by your organization. This process has five steps:

1. Create a private certificate authority (CA) using Amazon Private CA
2. Issue and install the root certificate on the private CA
3. Request a private certificate from Amazon Certificate Manager for your file system and clients
4. Export the certificate for the file system and clients.

For more information, see [Private CA administration](#) in the Amazon Private Certificate Authority User Guide.

To create the root private CA

1. When you create a CA, you must specify the CA configuration in a file that you supply. The following command uses the Nano text editor to create the `ca_config.txt` file, which specifies the following information:
 - The name of the algorithm
 - The signing algorithm that the CA uses to sign
 - X.500 subject information

```
$ > nano ca_config.txt
```

The text editor appears.

2. Edit the file with the specifications for your CA.

```
{
  "KeyAlgorithm":"RSA_2048",
  "SigningAlgorithm":"SHA256WITHRSA",
  "Subject":{
    "Country":"US",
    "Organization":"Example Corp",
    "OrganizationalUnit":"Sales",
    "State":"WA",
```

```

    "Locality": "Seattle",
    "CommonName": "*.ec2.internal"
  }
}

```

3. Save and close the file, exiting the text editor. For more information, see [Procedure for creating a CA](#) in the Amazon Private Certificate Authority User Guide.
4. Use the [create-certificate-authority](#) Amazon Private CA CLI command to create a private CA.

```

~/home > aws acm-pca create-certificate-authority \
  --certificate-authority-configuration file://ca_config.txt \
  --certificate-authority-type "ROOT" \
  --idempotency-token 01234567 --region aws-region

```

If successful, this command outputs the Amazon Resource Name (ARN) of the CA.

```

{
  "CertificateAuthorityArn": "arn:aws:acm-pca:aws-region:111122223333:certificate-
authority/12345678-1234-1234-1234-123456789012"
}

```

To create and install a certificate for your private root CA (Amazon CLI)

1. Generate a certificate signing request (CSR) using the [get-certificate-authority-csr](#) Amazon CLI command.

```

$ aws acm-pca get-certificate-authority-csr \
  --certificate-authority-arn arn:aws:acm-pca:aws-
region:111122223333:certificate-authority/12345678-1234-1234-1234-123456789012 \
  --output text \
  --endpoint https://acm-pca.aws-region.amazonaws.com \
  --region eu-west-1 > ca.csr

```

The resulting file `ca.csr`, a PEM file encoded in base64 format, has the following appearance.

```

-----BEGIN CERTIFICATE-----
MIICiTCCAfICQD6m7oRw0uX0jANBgkqhkiG9w0BAQUFADCBiDELMAkGA1UEBhMC
VVMxCzAJBgNVBAgTAldBMRAdDgYDVQQHEwdTZWF0dGx1MQ8wDQYDVQQKEwZBbWFG
b24xFDASBgNVBAScT01BTSBDb25zb2x1MRIwEAYDVQQDEw1UZXRhbnQ21sYWMxHzAd
BgkqhkiG9w0BCQEWEG5vb25lQGFtYXN0LmVudG9wHhcNMTEwNDI1MjA0NTIxWhcN

```

```
MTIwNDI0MjA0NTIxWjCBiDELMAKGA1UEBhMCMVVMxCzAJBgNVBAGTA1dBMRwDgYD
VQQHEwdTZWF0dGx1MQ8wDQYDVQQKEwZBbWF6b24xZDASBgNVBAwTC0lBTSBDb25z
b2x1MRIwEAYDVQQDEw1UZXR0Q21sYWxhZAdBgkqhkiG9w0BCQEWEG5vb251QGft
YXpvbi5jb20wZGZ8wDQYJKoZIhvcNAQEBBQADgY0AMIGJAoGBAMaK0dn+a4GmWIWJ
21uUSfwfEvySwTc2XADZ4nB+BLYgVIk60CpiwsZ3G93vUEI03IyNoH/f0wYK8m9T
rDHudUZg3qX4waLG5M43q7Wgc/MbQITx0USQv7c7ugFFDzQGBzZswY6786m86gpE
Ibb30hjZncvQAaRHhd1QWIMm2nrAgMBAAEwDQYJKoZIhvcNAQEFBQADgYEAtCu4
nUHVvXyUntneD9+h8Mg9q6q+auNKyExzyLwaxlAoo7TJHidbtS4J5iNmZgXL0Fkb
FFBjvSfpJI1J00zbhNYS5f6GuoEDmFJl0ZxBHjJnyp3780D8uTs7fLvJx79LjStb
NYiytVbZPQUQ5Yaxu2jXnimvw3rrszlaEXAMPLE=
-----END CERTIFICATE-----
```

For more information, see [Installing a root CA certificate](#) in the Amazon Private Certificate Authority User Guide.

2. Use the [issue-certificate](#) Amazon CLI command to issue and install the root certificate on your private CA.

```
$ aws acm-pca issue-certificate \
  --certificate-authority-arn arn:aws:acm-pca:aws-
  region:111122223333:certificate-authority/12345678-1234-1234-1234-123456789012 \
  --csr file://ca.csr \
  --signing-algorithm SHA256WITHRSA \
  --template-arn arn:aws:acm-pca::template/RootCACertificate/V1 \
  --validity Value=3650,Type=DAYS --region aws-region
```

3. Download the root certificate using the [get-certificate](#) Amazon CLI command.

```
$ aws acm-pca get-certificate \
  --certificate-authority-arn arn:aws:acm-pca:aws-
  region:111122223333:certificate-authority/12345678-1234-1234-1234-123456789012 \
  --certificate-arn arn:aws:acm-pca:aws-region:486768734100:certificate-
  authority/12345678-1234-1234-1234-123456789012/certificate/
  abcdef0123456789abcdef0123456789 \
  --output text --region aws-region > rootCA.pem
```

4. Install the root certificate on your private CA using the [import-certificate-authority-certificate](#) Amazon CLI command.

```
$ aws acm-pca import-certificate-authority-certificate \
  --certificate-authority-arn arn:aws:acm-pca:aws-
  region:111122223333:certificate-authority/12345678-1234-1234-1234-123456789012 \
```

```
--certificate file://rootCA.pem --region aws-region
```

Generate and export the file system and client certificate

1. Use the [request-certificate](#) Amazon CLI command to request an Amazon Certificate Manager certificate to use on your file system and clients.

```
$ aws acm request-certificate \
  --domain-name *.ec2.internal \
  --idempotency-token 12345 \
  --region aws-region \
  --certificate-authority-arn arn:aws:acm-pca:aws-
region:111122223333:certificate-authority/12345678-1234-1234-1234-123456789012
```

If the request is successful, the ARN of the issued certificate is returned.

2. For security, you must assign a passphrase for the private key when exporting it. Create a passphrase and store it in a file named `passphrase.txt`
3. Use the [export-certificate](#) Amazon CLI command to export the private certificate issued previously. The exported file contains the certificate, the certificate chain, and the encrypted private 2048-bit RSA key associated with the public key that is embedded in the certificate. For security, you must assign a passphrase for the private key when exporting it. The following example is for a Linux EC2 instance.

```
$ aws acm export-certificate \
  --certificate-arn arn:aws:acm:aws-
region:111122223333:certificate/12345678-1234-1234-1234-123456789012 \
  --passphrase $(cat passphrase.txt | base64) --region aws-region >
  exported_cert.json
```

4. Use the following `jq` commands to extract the private key and the certificate from the JSON response.

```
$ cat exported_cert.json | jq -r .PrivateKey > prv.key

cat exported_cert.json | jq -r .Certificate > cert.pem
openssl rsa -in prv.key -passin pass:$passphrase -out decrypted.key
```

5. Use the following `openssl` command to decrypt the private key from the JSON response. After entering the command, you are prompted for the passphrase.

```
$ openssl rsa -in prv.key -passin pass:$passphrase -out decrypted.key
```

Installing and configuring Libreswan IPsec on an Amazon Linux 2 client

The following sections provide instructions for installing and configuring Libreswan IPsec on an Amazon EC2 instance running Amazon Linux 2.

To install and configure Libreswan

1. Connect to your EC2 instance using SSH. For specific instructions on how to do this, see [Connect to your Linux instance using an SSH client](#) in the Amazon Elastic Compute Cloud User Guide for Linux Instances.
2. Run the following command to install libreswan:

```
$ sudo yum install libreswan
```

3. (Optional) When verifying IPsec in a later step, these properties might be flagged without these settings. We suggest testing your set up first without these settings. If your connection has problems, return to this step and make the following changes.

After the installation completes, use your preferred text editor to add the following entries to the `/etc/sysctl.conf` file.

```
net.ipv4.ip_forward=1
net.ipv4.conf.all.accept_redirects = 0
net.ipv4.conf.all.secure_redirects = 0
net.ipv4.conf.all.send_redirects = 0
net.ipv4.conf.default.accept_redirects = 0
net.ipv4.conf.default.send_redirects = 0
net.ipv4.conf.lo.accept_redirects = 0
net.ipv4.conf.lo.send_redirects = 0
net.ipv4.conf.all.rp_filter = 0
net.ipv4.conf.default.rp_filter = 0
net.ipv4.conf.eth0.rp_filter = 0
```

Save the changes and exit the text editor.

4. Apply the changes.

```
$ sudo sysctl -p
```

5. Verify the IPsec configuration.

```
$ sudo ipsec verify
```

Verify that the version of Libreswan you installed is running.

6. Initialize the IPsec NSS database.

```
$ sudo ipsec checknss
```

To install the certificate on the client

1. Copy the [certificate you generated](#) for the client to the working directory on the EC2 instance.
You
2. Export the certificate generated previously into a format compatible with libreswan.

```
$ openssl pkcs12 -export -in cert.pem -inkey decrypted.key \  
-certfile rootCA.pem -out certkey.p12 -name fsx
```

3. Import the reformatted key, providing the passphrase when prompted.

```
$ sudo ipsec import certkey.p12
```

4. Create an IPsec configuration file using the preferred text editor.

```
$ sudo cat /etc/ipsec.d/nfs.conf
```

Add the following entries to the config file:

```
conn fsxn  
  authby=rsasig  
  left=172.31.77.6  
  right=198.19.254.13  
  auto=start  
  type=transport  
  ikev2=insist  
  keyexchange=ike
```

```
ike=aes256-sha2_384;dh20
esp=aes_gcm_c256
leftcert=fsx
lefttrsasigkey=%cert
leftid=%fromcert
rightid=%fromcert
righttrsasigkey=%cert
```

You will start IPsec on the client after configuring IPsec on your file system.

Configuring IPsec on your file system

This section provides instructions on installing the certificate on your FSx for ONTAP file system, and configuring IPsec.

To install the certificate on your file system

1. Copy the root certificate (`rootCA.pem`), the client certificate (`cert.pem`) and the decrypted key (`decrypted.key`) files to your file system. You will need to know the passphrase for the certificate.
2. To access the NetApp ONTAP CLI, establish an SSH session on the management port of the Amazon FSx for NetApp ONTAP file system by running the following command. Replace *management_endpoint_ip* with the IP address of the file system's management port.

```
[~]$ ssh fsxadmin@management_endpoint_ip
```

For more information, see [Managing file systems with the ONTAP CLI](#).

3. Use `cat` on a client (not on your file system) to list out the contents of the `rootCA.pem`, `cert.pem` and `decrypted.key` files so that you can copy the output of each file and paste it when prompted in the following steps.

```
$ > cat cert.pem
```

Copy the certificate contents.

4. You must install all CA certificates used during the mutual authentication, including both ONTAP-side and client-side CAs, to ONTAP certificate management unless it is already installed (as is the case of an ONTAP self-signed root-CA).

Use the `security certificate install` NetApp CLI command as follows to install the client certificate:

```
FSxID123:: > security certificate install -vserver dr -type client -cert-name ipsec-client-cert
```

Please enter Certificate: Press <Enter> when done

Paste in the contents of the `cert.pem` file that you copied previously and press Enter.

Please enter Private Key: Press <Enter> when done

Paste in the contents of the `decrypted.key` file, and press enter.

Do you want to continue entering root and/or intermediate certificates {y|n}:

Enter `n` to complete entering the client certificate.

5. Create and install a certificate for use by the SVM. The issuer CA of this certificate must already be installed to ONTAP and added to IPsec.

Use the following command to install the root certificate.

```
FSxID123:: > security certificate install -vserver dr -type server-ca -cert-name ipsec-ca-cert
```

Please enter Certificate: Press <Enter> when done

Paste in the contents of the `rootCA.pem` file, and press enter.

6. To ensure that the CA installed is within the IPsec CA searching path during authentication, add the ONTAP certificate management CAs to the IPsec module using the “`security ipsec ca-certificate add`” command.

Enter the following command to add the root certificate.

```
FSxID123:: > security ipsec ca-certificate add -vserver dr -ca-certs ipsec-ca-cert
```


7. Enter the following command to create the required IPsec policy in the security policy database (SPD).

```
security ipsec policy create -vserver dr -name policy-name -local-ip-  
subnets 198.19.254.13/32 -remote-ip-subnets 172.31.0.0/16 -auth-method PKI -action  
ESP_TRA -cipher-suite SUITEB_GCM256 -cert-name ipsec-client-cert -local-identity  
"CN=*.ec2.internal" -remote-identity "CN=*.ec2.internal"
```

8. Use the following command to show the IPsec policy for the file system to confirm.

```
FSxID123:: > security ipsec policy show -vserver dr -instance
```

```

                Vserver: dr
                Policy Name: promise
                Local IP Subnets: 198.19.254.13/32
                Remote IP Subnets: 172.31.0.0/16
                Local Ports: 0-0
                Remote Ports: 0-0
                Protocols: any
                Action: ESP_TRA
                Cipher Suite: SUITEB_GCM256
                IKE Security Association Lifetime: 86400
                IPsec Security Association Lifetime: 28800
                IPsec Security Association Lifetime (bytes): 0
                Is Policy Enabled: true
                Local Identity: CN=*.ec2.internal
                Remote Identity: CN=*.ec2.internal
                Authentication Method: PKI
                Certificate for Local Identity: ipsec-client-cert
```

Start IPsec on the client

Now IPsec is configured on both the FSx for ONTAP file system and the client, you can start IPsec on the client.

1. Connect to your client system using SSH.
2. Start IPsec.

```
$ sudo ipsec start
```

3. Check the status of IPsec.

```
$ sudo ipsec status
```

4. Mount a volume on your file system.

```
$ sudo mount -t nfs 198.19.254.13:/benchmark /home/ec2-user/acm/dr
```

5. Verify the IPsec setup by showing the encrypted connection on your FSx for ONTAP file system.

```
FSxID123:: > security ipsec show-ikesa -node FsxId123
FsxId08ac16c7ec2781a58::> security ipsec show-ikesa -node FsxId08ac16c7ec2781a58-01
```

Vserver	Policy Name	Local Address	Remote Address	Initiator-SPI	State
dr	<i>policy-name</i>	198.19.254.13	172.31.77.6	551c55de57fe8976	ESTABLISHED
fsx	<i>policy-name</i>	198.19.254.38	172.31.65.193	4fd3f22c993e60c5	ESTABLISHED

2 entries were displayed.

Setting up IPsec for multiple clients

When a small number of clients need to leverage IPsec, using a single SPD entry for each client is sufficient. However, when hundreds or even thousands of clients need to leverage IPsec, we recommend that you use IPsec multiple client configuration.

FSx for ONTAP supports connecting multiple clients across many networks to a single SVM IP address with IPsec enabled. You can accomplish this using either the subnet configuration or the Allow all clients configuration, which are explained in the following procedures:

To configure IPsec for multiple clients using a subnet configuration

To allow all clients on a particular subnet (192.168.134.0/24 for example) to connect to a single SVM IP address using a single SPD policy entry, you must specify the `remote-ip-subnets` in subnet form. Additionally, you must specify the `remote-identity` field with the correct client side identity.

⚠ Important

When using certificate authentication, each client can use either their own unique certificate or a shared certificate to authenticate. FSx for ONTAP IPsec checks the validity of the certificate based on the CAs installed on its local trust store. FSx for ONTAP also supports certificate revocation list (CRL) checking.

1. To access the NetApp ONTAP CLI, establish an SSH session on the management port of the Amazon FSx for NetApp ONTAP file system by running the following command. Replace *management_endpoint_ip* with the IP address of the file system's management port.

```
[~]$ ssh fsxadmin@management_endpoint_ip
```

For more information, see [Managing file systems with the ONTAP CLI](#).

2. Use the security ipsec policy create NetApp ONTAP CLI command as follows, replacing the *sample* values with your specific values.

```
FsxId123456::> security ipsec policy create -vserver svm_name -name policy_name \  
-local-ip-subnets 192.168.134.34/32 -remote-ip-subnets 192.168.134.0/24 \  
-local-ports 2049 -protocols tcp -auth-method PSK \  
-cert-name my_nfs_server_cert -local-identity ontap_side_identity \  
-remote-identity client_side_identity
```

To configure IPsec for multiple clients using an allow all clients configuration

To allow any client, regardless of their source IP address, to connect to the SVM IPsec-enabled IP address, use the *0.0.0.0/0* wild card when specifying the `remote-ip-subnets` field.

Additionally, you must specify the `remote-identity` field with the correct client side identity. For certificate authentication, you can enter ANYTHING.

Also, when the *0.0.0.0/0* wild card is used, you must configure a specific local or remote port number to use. For example, NFS port 2049.

1. To access the NetApp ONTAP CLI, establish an SSH session on the management port of the Amazon FSx for NetApp ONTAP file system by running the following command. Replace *management_endpoint_ip* with the IP address of the file system's management port.

```
[~]$ ssh fsxadmin@management_endpoint_ip
```

For more information, see [Managing file systems with the ONTAP CLI](#).

2. Use the security ipsec policy create NetApp ONTAP CLI command as follows, replacing the *sample* values with your specific values.

```
FsxId123456::> security ipsec policy create -vserver svm_name -name policy_name \  
-local-ip-subnets 192.168.134.34/32 -remote-ip-subnets 0.0.0.0/0 \  
-local-ports 2049 -protocols tcp -auth-method PSK \  
-cert-name my_nfs_server_cert -local-identity ontap_side_identity \  
-local-ports 2049 -remote-identity client_side_identity
```

Identity and access management for Amazon FSx for NetApp ONTAP

Amazon Identity and Access Management (IAM) is an Amazon Web Service that helps an administrator securely control access to Amazon resources. IAM administrators control who can be *authenticated* (signed in) and *authorized* (have permissions) to use Amazon FSx resources. IAM is an Amazon Web Service that you can use with no additional charge.

Topics

- [Audience](#)
- [Authenticating with identities](#)
- [Managing access using policies](#)
- [How Amazon FSx for NetApp ONTAP works with IAM](#)
- [Identity-based policy examples for Amazon FSx for NetApp ONTAP](#)
- [Troubleshooting Amazon FSx for NetApp ONTAP identity and access](#)
- [Using tags with Amazon FSx](#)
- [Using service-linked roles for Amazon FSx](#)

Audience

How you use Amazon Identity and Access Management (IAM) differs, depending on the work that you do in Amazon FSx.

Service user – If you use the Amazon FSx service to do your job, then your administrator provides you with the credentials and permissions that you need. As you use more Amazon FSx features to do your work, you might need additional permissions. Understanding how access is managed can help you request the right permissions from your administrator. If you cannot access a feature in Amazon FSx, see [Troubleshooting Amazon FSx for NetApp ONTAP identity and access](#).

Service administrator – If you're in charge of Amazon FSx resources at your company, you probably have full access to Amazon FSx. It's your job to determine which Amazon FSx features and resources your service users should access. You must then submit requests to your IAM administrator to change the permissions of your service users. Review the information on this page to understand the basic concepts of IAM. To learn more about how your company can use IAM with Amazon FSx, see [How Amazon FSx for NetApp ONTAP works with IAM](#).

IAM administrator – If you're an IAM administrator, you might want to learn details about how you can write policies to manage access to Amazon FSx. To view example Amazon FSx identity-based policies that you can use in IAM, see [Identity-based policy examples for Amazon FSx for NetApp ONTAP](#).

Authenticating with identities

Authentication is how you sign in to Amazon using your identity credentials. You must be *authenticated* (signed in to Amazon) as the Amazon Web Services account root user, as an IAM user, or by assuming an IAM role.

If you access Amazon programmatically, Amazon provides a software development kit (SDK) and a command line interface (CLI) to cryptographically sign your requests by using your credentials. If you don't use Amazon tools, you must sign requests yourself. For more information about using the recommended method to sign requests yourself, see [Signing Amazon API requests](#) in the *IAM User Guide*.

Regardless of the authentication method that you use, you might be required to provide additional security information. For example, Amazon recommends that you use multi-factor authentication (MFA) to increase the security of your account. To learn more, see [Using multi-factor authentication \(MFA\) in Amazon](#) in the *IAM User Guide*.

Amazon Web Services account root user

When you create an Amazon Web Services account, you begin with one sign-in identity that has complete access to all Amazon Web Services and resources in the account. This identity is called the Amazon Web Services account *root user* and is accessed by signing in with the email address and password that you used to create the account. We strongly recommend that you don't use the root user for your everyday tasks. Safeguard your root user credentials and use them to perform the tasks that only the root user can perform. For the complete list of tasks that require you to sign in as the root user, see [Tasks that require root user credentials](#) in the *IAM User Guide*.

Federated identity

As a best practice, require human users, including users that require administrator access, to use federation with an identity provider to access Amazon Web Services by using temporary credentials.

A *federated identity* is a user from your enterprise user directory, a web identity provider, the Amazon Directory Service, or any user that accesses Amazon Web Services by using credentials provided through an identity source. When federated identities access Amazon Web Services accounts, they assume roles, and the roles provide temporary credentials.

IAM users and groups

An [IAM user](#) is an identity within your Amazon Web Services account that has specific permissions for a single person or application. Where possible, we recommend relying on temporary credentials instead of creating IAM users who have long-term credentials such as passwords and access keys. However, if you have specific use cases that require long-term credentials with IAM users, we recommend that you rotate access keys. For more information, see [Rotate access keys regularly for use cases that require long-term credentials](#) in the *IAM User Guide*.

An [IAM group](#) is an identity that specifies a collection of IAM users. You can't sign in as a group. You can use groups to specify permissions for multiple users at a time. Groups make permissions easier to manage for large sets of users. For example, you could have a group named *IAMAdmins* and give that group permissions to administer IAM resources.

Users are different from roles. A user is uniquely associated with one person or application, but a role is intended to be assumable by anyone who needs it. Users have permanent long-term credentials, but roles provide temporary credentials. To learn more, see [When to create an IAM user \(instead of a role\)](#) in the *IAM User Guide*.

IAM roles

An [IAM role](#) is an identity within your Amazon Web Services account that has specific permissions. It is similar to an IAM user, but is not associated with a specific person. You can temporarily assume an IAM role in the Amazon Web Services Management Console by [switching roles](#). You can assume a role by calling an Amazon CLI or Amazon API operation or by using a custom URL. For more information about methods for using roles, see [Using IAM roles](#) in the *IAM User Guide*.

IAM roles with temporary credentials are useful in the following situations:

- **Federated user access** – To assign permissions to a federated identity, you create a role and define permissions for the role. When a federated identity authenticates, the identity is associated with the role and is granted the permissions that are defined by the role. For information about roles for federation, see [Creating a role for a third-party Identity Provider](#) in the *IAM User Guide*.
- **Temporary IAM user permissions** – An IAM user or role can assume an IAM role to temporarily take on different permissions for a specific task.
- **Cross-account access** – You can use an IAM role to allow someone (a trusted principal) in a different account to access resources in your account. Roles are the primary way to grant cross-account access. However, with some Amazon Web Services, you can attach a policy directly to a resource (instead of using a role as a proxy). To learn the difference between roles and resource-based policies for cross-account access, see [How IAM roles differ from resource-based policies](#) in the *IAM User Guide*.
- **Cross-service access** – Some Amazon Web Services use features in other Amazon Web Services. For example, when you make a call in a service, it's common for that service to run applications in Amazon EC2 or store objects in Amazon S3. A service might do this using the calling principal's permissions, using a service role, or using a service-linked role.
- **Forward access sessions (FAS)** – When you use an IAM user or role to perform actions in Amazon, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an Amazon Web Service, combined with the requesting Amazon Web Service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other Amazon Web Services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see [Forward access sessions](#).
- **Service role** – A service role is an [IAM role](#) that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For

more information, see [Creating a role to delegate permissions to an Amazon Web Service](#) in the *IAM User Guide*.

- **Service-linked role** – A service-linked role is a type of service role that is linked to an Amazon Web Service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your Amazon Web Services account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.
- **Applications running on Amazon EC2** – You can use an IAM role to manage temporary credentials for applications that are running on an EC2 instance and making Amazon CLI or Amazon API requests. This is preferable to storing access keys within the EC2 instance. To assign an Amazon role to an EC2 instance and make it available to all of its applications, you create an instance profile that is attached to the instance. An instance profile contains the role and enables programs that are running on the EC2 instance to get temporary credentials. For more information, see [Using an IAM role to grant permissions to applications running on Amazon EC2 instances](#) in the *IAM User Guide*.

To learn whether to use IAM roles or IAM users, see [When to create an IAM role \(instead of a user\)](#) in the *IAM User Guide*.

Managing access using policies

You control access in Amazon by creating policies and attaching them to Amazon identities or resources. A policy is an object in Amazon that, when associated with an identity or resource, defines their permissions. Amazon evaluates these policies when a principal (user, root user, or role session) makes a request. Permissions in the policies determine whether the request is allowed or denied. Most policies are stored in Amazon as JSON documents. For more information about the structure and contents of JSON policy documents, see [Overview of JSON policies](#) in the *IAM User Guide*.

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

By default, users and roles have no permissions. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

IAM policies define permissions for an action regardless of the method that you use to perform the operation. For example, suppose that you have a policy that allows the `iam:GetRole` action.

A user with that policy can get role information from the Amazon Web Services Management Console, the Amazon CLI, or the Amazon API.

Identity-based policies

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see [Creating IAM policies](#) in the *IAM User Guide*.

Identity-based policies can be further categorized as *inline policies* or *managed policies*. Inline policies are embedded directly into a single user, group, or role. Managed policies are standalone policies that you can attach to multiple users, groups, and roles in your Amazon Web Services account. Managed policies include Amazon managed policies and customer managed policies. To learn how to choose between a managed policy or an inline policy, see [Choosing between managed policies and inline policies](#) in the *IAM User Guide*.

Resource-based policies

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must [specify a principal](#) in a resource-based policy. Principals can include accounts, users, roles, federated users, or Amazon Web Services.

Resource-based policies are inline policies that are located in that service. You can't use Amazon managed policies from IAM in a resource-based policy.

Access control lists (ACLs)

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Amazon S3, Amazon WAF, and Amazon VPC are examples of services that support ACLs. To learn more about ACLs, see [Access control list \(ACL\) overview](#) in the *Amazon Simple Storage Service Developer Guide*.

Other policy types

Amazon supports additional, less-common policy types. These policy types can set the maximum permissions granted to you by the more common policy types.

- **Permissions boundaries** – A permissions boundary is an advanced feature in which you set the maximum permissions that an identity-based policy can grant to an IAM entity (IAM user or role). You can set a permissions boundary for an entity. The resulting permissions are the intersection of an entity's identity-based policies and its permissions boundaries. Resource-based policies that specify the user or role in the `Principal` field are not limited by the permissions boundary. An explicit deny in any of these policies overrides the allow. For more information about permissions boundaries, see [Permissions boundaries for IAM entities](#) in the *IAM User Guide*.
- **Service control policies (SCPs)** – SCPs are JSON policies that specify the maximum permissions for an organization or organizational unit (OU) in Amazon Organizations. Amazon Organizations is a service for grouping and centrally managing multiple Amazon Web Services accounts that your business owns. If you enable all features in an organization, then you can apply service control policies (SCPs) to any or all of your accounts. The SCP limits permissions for entities in member accounts, including each Amazon Web Services account root user. For more information about Organizations and SCPs, see [How SCPs work](#) in the *Amazon Organizations User Guide*.
- **Session policies** – Session policies are advanced policies that you pass as a parameter when you programmatically create a temporary session for a role or federated user. The resulting session's permissions are the intersection of the user or role's identity-based policies and the session policies. Permissions can also come from a resource-based policy. An explicit deny in any of these policies overrides the allow. For more information, see [Session policies](#) in the *IAM User Guide*.

Multiple policy types

When multiple types of policies apply to a request, the resulting permissions are more complicated to understand. To learn how Amazon determines whether to allow a request when multiple policy types are involved, see [Policy evaluation logic](#) in the *IAM User Guide*.

How Amazon FSx for NetApp ONTAP works with IAM

Before you use IAM to manage access to Amazon FSx, learn what IAM features are available to use with Amazon FSx.

IAM features you can use with Amazon FSx for NetApp ONTAP

IAM feature	Amazon FSx support
Identity-based policies	Yes
Resource-based policies	No
Policy actions	Yes
Policy resources	Yes
Policy condition keys	Yes
ACLs	No
ABAC (tags in policies)	Yes
Temporary credentials	Yes
Forward access sessions (FAS)	Yes
Service roles	No
Service-linked roles	Yes

To get a high-level view of how Amazon FSx and other Amazon services work with most IAM features, see [Amazon services that work with IAM](#) in the *IAM User Guide*.

Identity-based policies for Amazon FSx

Supports identity-based policies	Yes
----------------------------------	-----

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see [Creating IAM policies](#) in the *IAM User Guide*.

With IAM identity-based policies, you can specify allowed or denied actions and resources as well as the conditions under which actions are allowed or denied. You can't specify the principal in an identity-based policy because it applies to the user or role to which it is attached. To learn about all of the elements that you can use in a JSON policy, see [IAM JSON policy elements reference](#) in the *IAM User Guide*.

Identity-based policy examples for Amazon FSx

To view examples of Amazon FSx identity-based policies, see [Identity-based policy examples for Amazon FSx for NetApp ONTAP](#).

Resource-based policies within Amazon FSx

Supports resource-based policies	No
----------------------------------	----

Policy actions for Amazon FSx

Supports policy actions	Yes
-------------------------	-----

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Action` element of a JSON policy describes the actions that you can use to allow or deny access in a policy. Policy actions usually have the same name as the associated Amazon API operation. There are some exceptions, such as *permission-only actions* that don't have a matching API operation. There are also some operations that require multiple actions in a policy. These additional actions are called *dependent actions*.

Include actions in a policy to grant permissions to perform the associated operation.

To see a list of Amazon FSx actions, see [Actions defined by Amazon FSx](#) in the *Service Authorization Reference*.

Policy actions in Amazon FSx use the following prefix before the action:

```
fsx
```

To specify multiple actions in a single statement, separate them with commas.

```
"Action": [  
  "fsx:action1",  
  "fsx:action2"  
]
```

To view examples of Amazon FSx identity-based policies, see [Identity-based policy examples for Amazon FSx for NetApp ONTAP](#).

Policy resources for Amazon FSx

Supports policy resources

Yes

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Resource JSON policy element specifies the object or objects to which the action applies. Statements must include either a Resource or a NotResource element. As a best practice, specify a resource using its [Amazon Resource Name \(ARN\)](#). You can do this for actions that support a specific resource type, known as *resource-level permissions*.

For actions that don't support resource-level permissions, such as listing operations, use a wildcard (*) to indicate that the statement applies to all resources.

```
"Resource": "*"
```

To see a list of Amazon FSx resource types and their ARNs, see [Resources defined by Amazon FSx](#) in the *Service Authorization Reference*. To learn with which actions you can specify the ARN of each resource, see [Actions defined by Amazon FSx](#).

To view examples of Amazon FSx identity-based policies, see [Identity-based policy examples for Amazon FSx for NetApp ONTAP](#).

Policy condition keys for Amazon FSx

Supports service-specific policy condition keys Yes

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Condition` element (or *Condition block*) lets you specify conditions in which a statement is in effect. The `Condition` element is optional. You can create conditional expressions that use [condition operators](#), such as equals or less than, to match the condition in the policy with values in the request.

If you specify multiple `Condition` elements in a statement, or multiple keys in a single `Condition` element, Amazon evaluates them using a logical AND operation. If you specify multiple values for a single condition key, Amazon evaluates the condition using a logical OR operation. All of the conditions must be met before the statement's permissions are granted.

You can also use placeholder variables when you specify conditions. For example, you can grant an IAM user permission to access a resource only if it is tagged with their IAM user name. For more information, see [IAM policy elements: variables and tags](#) in the *IAM User Guide*.

Amazon supports global condition keys and service-specific condition keys. To see all Amazon global condition keys, see [Amazon global condition context keys](#) in the *IAM User Guide*.

To see a list of Amazon FSx condition keys, see [Condition keys for Amazon FSx](#) in the *Service Authorization Reference*. To learn with which actions and resources you can use a condition key, see [Actions defined by Amazon FSx](#).

To view examples of Amazon FSx identity-based policies, see [Identity-based policy examples for Amazon FSx for NetApp ONTAP](#).

Access control lists (ACLs) in Amazon FSx

Supports ACLs No

Attribute-based access control (ABAC) with Amazon FSx

Supports ABAC (tags in policies)	Yes
----------------------------------	-----

Attribute-based access control (ABAC) is an authorization strategy that defines permissions based on attributes. In Amazon, these attributes are called *tags*. You can attach tags to IAM entities (users or roles) and to many Amazon resources. Tagging entities and resources is the first step of ABAC. Then you design ABAC policies to allow operations when the principal's tag matches the tag on the resource that they are trying to access.

ABAC is helpful in environments that are growing rapidly and helps with situations where policy management becomes cumbersome.

To control access based on tags, you provide tag information in the [condition element](#) of a policy using the `aws:ResourceTag/key-name`, `aws:RequestTag/key-name`, or `aws:TagKeys` condition keys.

If a service supports all three condition keys for every resource type, then the value is **Yes** for the service. If a service supports all three condition keys for only some resource types, then the value is **Partial**.

For more information about ABAC, see [What is ABAC?](#) in the *IAM User Guide*. To view a tutorial with steps for setting up ABAC, see [Use attribute-based access control \(ABAC\)](#) in the *IAM User Guide*.

For more information about tagging Amazon FSx resources, see [Tag your Amazon FSx resources](#).

To view an example identity-based policy for limiting access to a resource based on the tags on that resource, see [Using tags to control access to your Amazon FSx resources](#).

Using Temporary credentials with Amazon FSx

Supports temporary credentials	Yes
--------------------------------	-----

Some Amazon Web Services don't work when you sign in using temporary credentials. For additional information, including which Amazon Web Services work with temporary credentials, see [Amazon Web Services that work with IAM](#) in the *IAM User Guide*.

You are using temporary credentials if you sign in to the Amazon Web Services Management Console using any method except a user name and password. For example, when you access Amazon using your company's single sign-on (SSO) link, that process automatically creates temporary credentials. You also automatically create temporary credentials when you sign in to the console as a user and then switch roles. For more information about switching roles, see [Switching to a role \(console\)](#) in the *IAM User Guide*.

You can manually create temporary credentials using the Amazon CLI or Amazon API. You can then use those temporary credentials to access Amazon. Amazon recommends that you dynamically generate temporary credentials instead of using long-term access keys. For more information, see [Temporary security credentials in IAM](#).

Forward access sessions for Amazon FSx

Supports forward access sessions (FAS)	Yes
--	-----

When you use an IAM user or role to perform actions in Amazon, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an Amazon Web Service, combined with the requesting Amazon Web Service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other Amazon Web Services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see [Forward access sessions](#).

Service roles for Amazon FSx

Supports service roles	No
------------------------	----

Service-linked roles for Amazon FSx

Supports service-linked roles	Yes
-------------------------------	-----

A service-linked role is a type of service role that is linked to an Amazon Web Service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your

Amazon Web Services account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.

For details about creating or managing Amazon FSx service-linked roles, see [Using service-linked roles for Amazon FSx](#).

Identity-based policy examples for Amazon FSx for NetApp ONTAP

By default, users and roles don't have permission to create or modify Amazon FSx resources. They also can't perform tasks by using the Amazon Web Services Management Console, Amazon Command Line Interface (Amazon CLI), or Amazon API. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

To learn how to create an IAM identity-based policy by using these example JSON policy documents, see [Creating IAM policies](#) in the *IAM User Guide*.

For details about actions and resource types defined by Amazon FSx, including the format of the ARNs for each of the resource types, see [Actions, resources, and condition keys for Amazon FSx](#) in the *Service Authorization Reference*.

Topics

- [Policy best practices](#)
- [Using the Amazon FSx console](#)
- [Allow users to view their own permissions](#)

Policy best practices

Identity-based policies determine whether someone can create, access, or delete Amazon FSx resources in your account. These actions can incur costs for your Amazon Web Services account. When you create or edit identity-based policies, follow these guidelines and recommendations:

- **Get started with Amazon managed policies and move toward least-privilege permissions**
 - To get started granting permissions to your users and workloads, use the *Amazon managed policies* that grant permissions for many common use cases. They are available in your Amazon Web Services account. We recommend that you reduce permissions further by defining Amazon customer managed policies that are specific to your use cases. For more information, see [Amazon managed policies](#) or [Amazon managed policies for job functions](#) in the *IAM User Guide*.

- **Apply least-privilege permissions** – When you set permissions with IAM policies, grant only the permissions required to perform a task. You do this by defining the actions that can be taken on specific resources under specific conditions, also known as *least-privilege permissions*. For more information about using IAM to apply permissions, see [Policies and permissions in IAM](#) in the *IAM User Guide*.
- **Use conditions in IAM policies to further restrict access** – You can add a condition to your policies to limit access to actions and resources. For example, you can write a policy condition to specify that all requests must be sent using SSL. You can also use conditions to grant access to service actions if they are used through a specific Amazon Web Service, such as Amazon CloudFormation. For more information, see [IAM JSON policy elements: Condition](#) in the *IAM User Guide*.
- **Use IAM Access Analyzer to validate your IAM policies to ensure secure and functional permissions** – IAM Access Analyzer validates new and existing policies so that the policies adhere to the IAM policy language (JSON) and IAM best practices. IAM Access Analyzer provides more than 100 policy checks and actionable recommendations to help you author secure and functional policies. For more information, see [IAM Access Analyzer policy validation](#) in the *IAM User Guide*.
- **Require multi-factor authentication (MFA)** – If you have a scenario that requires IAM users or a root user in your Amazon Web Services account, turn on MFA for additional security. To require MFA when API operations are called, add MFA conditions to your policies. For more information, see [Configuring MFA-protected API access](#) in the *IAM User Guide*.

For more information about best practices in IAM, see [Security best practices in IAM](#) in the *IAM User Guide*.

Using the Amazon FSx console

To access the Amazon FSx for NetApp ONTAP console, you must have a minimum set of permissions. These permissions must allow you to list and view details about the Amazon FSx resources in your Amazon Web Services account. If you create an identity-based policy that is more restrictive than the minimum required permissions, the console won't function as intended for entities (users or roles) with that policy.

You don't need to allow minimum console permissions for users that are making calls only to the Amazon CLI or the Amazon API. Instead, allow access to only the actions that match the API operation that they're trying to perform.

To ensure that users and roles can still use the Amazon FSx console, also attach the `AmazonFSxConsoleReadOnlyAccess` Amazon managed policy to the entities. For more information, see [Adding permissions to a user](#) in the *IAM User Guide*.

You can see the `AmazonFSxConsoleReadOnlyAccess` and other Amazon FSx managed service policies in [Amazon managed policies for Amazon FSx](#).

Allow users to view their own permissions

This example shows how you might create a policy that allows IAM users to view the inline and managed policies that are attached to their user identity. This policy includes permissions to complete this action on the console or programmatically using the Amazon CLI or Amazon API.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ViewOwnUserInfo",
      "Effect": "Allow",
      "Action": [
        "iam:GetUserPolicy",
        "iam:ListGroupsWithUser",
        "iam:ListAttachedUserPolicies",
        "iam:ListUserPolicies",
        "iam:GetUser"
      ],
      "Resource": ["arn:aws-cn:iam::*:user/${aws:username}"]
    },
    {
      "Sid": "NavigateInConsole",
      "Effect": "Allow",
      "Action": [
        "iam:GetGroupPolicy",
        "iam:GetPolicyVersion",
        "iam:GetPolicy",
        "iam:ListAttachedGroupPolicies",
        "iam:ListGroupPolicies",
        "iam:ListPolicyVersions",
        "iam:ListPolicies",
        "iam:ListUsers"
      ],
      "Resource": "*"
    }
  ]
}
```

```
]
}
```

Troubleshooting Amazon FSx for NetApp ONTAP identity and access

Use the following information to help you diagnose and fix common issues that you might encounter when working with Amazon FSx and IAM.

Topics

- [I am not authorized to perform an action in Amazon FSx](#)
- [I am not authorized to perform iam:PassRole](#)
- [I want to allow people outside of my Amazon Web Services account to access my Amazon FSx resources](#)

I am not authorized to perform an action in Amazon FSx

If you receive an error that you're not authorized to perform an action, your policies must be updated to allow you to perform the action.

The following example error occurs when the mateojackson IAM user tries to use the console to view details about a fictional *my-example-widget* resource but doesn't have the fictional `fsx:GetWidget` permissions.

```
User: arn:aws-cn:iam::123456789012:user/mateojackson is not authorized to perform:
fsx:GetWidget on resource: my-example-widget
```

In this case, the policy for the mateojackson user must be updated to allow access to the *my-example-widget* resource by using the `fsx:GetWidget` action.

If you need help, contact your Amazon administrator. Your administrator is the person who provided you with your sign-in credentials.

I am not authorized to perform iam:PassRole

If you receive an error that you're not authorized to perform the `iam:PassRole` action, your policies must be updated to allow you to pass a role to Amazon FSx.

Some Amazon Web Services allow you to pass an existing role to that service instead of creating a new service role or service-linked role. To do this, you must have permissions to pass the role to the service.

The following example error occurs when an IAM user named `marymajor` tries to use the console to perform an action in Amazon FSx. However, the action requires the service to have permissions that are granted by a service role. Mary does not have permissions to pass the role to the service.

```
User: arn:aws-cn:iam::123456789012:user/marymajor is not authorized to perform:
iam:PassRole
```

In this case, Mary's policies must be updated to allow her to perform the `iam:PassRole` action.

If you need help, contact your Amazon administrator. Your administrator is the person who provided you with your sign-in credentials.

I want to allow people outside of my Amazon Web Services account to access my Amazon FSx resources

You can create a role that users in other accounts or people outside of your organization can use to access your resources. You can specify who is trusted to assume the role. For services that support resource-based policies or access control lists (ACLs), you can use those policies to grant people access to your resources.

To learn more, consult the following:

- To learn whether Amazon FSx supports these features, see [How Amazon FSx for NetApp ONTAP works with IAM](#).
- To learn how to provide access to your resources across Amazon Web Services accounts that you own, see [Providing access to an IAM user in another Amazon Web Services account that you own in the IAM User Guide](#).
- To learn how to provide access to your resources to third-party Amazon Web Services accounts, see [Providing access to Amazon Web Services accounts owned by third parties](#) in the *IAM User Guide*.
- To learn how to provide access through identity federation, see [Providing access to externally authenticated users \(identity federation\)](#) in the *IAM User Guide*.
- To learn the difference between using roles and resource-based policies for cross-account access, see [How IAM roles differ from resource-based policies](#) in the *IAM User Guide*.

Using tags with Amazon FSx

You can use tags to control access to Amazon FSx resources and to implement attribute-based access control (ABAC). To apply tags to Amazon FSx resources during creation, users must have certain Amazon Identity and Access Management (IAM) permissions.

Grant permission to tag resources during creation

With some resource-creating Amazon FSx API actions, you can specify tags when you create the resource. You can use these resource tags to implement attribute-based access control (ABAC). For more information, see [What is ABAC for Amazon?](#) in the *IAM User Guide*.

For users to tag resources on creation, they must have permission to use the action that creates the resource, such as `fsx:CreateFileSystem`, `fsx:CreateStorageVirtualMachine`, or `fsx:CreateVolume`. If tags are specified in the resource-creating action, IAM performs additional authorization on the `fsx:TagResource` action to verify if users have permissions to create tags. Therefore, users must also have explicit permissions to use the `fsx:TagResource` action.

The following example policy allows users to create file systems and storage virtual machines (SVMs) and apply tags to them during creation in a specific Amazon Web Services account.

```
{
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "fsx:CreateFileSystem",
        "fsx:CreateStorageVirtualMachine",
        "fsx:TagResource"
      ],
      "Resource": [
        "arn:aws:fsx:region:account-id:file-system/*",
        "arn:aws:fsx:region:account-id:file-system/*/storage-virtual-machine/*"
      ]
    }
  ]
}
```

Similarly, the following policy allows users to create backups on a specific file system and apply any tags to the backup during backup creation.

```
{
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "fsx:CreateBackup"
      ],
      "Resource": "arn:aws:fsx:region:account-id:file-system/file-system-id*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "fsx:TagResource"
      ],
      "Resource": "arn:aws:fsx:region:account-id:backup/*"
    }
  ]
}
```

The `fsx:TagResource` action is evaluated only if tags are applied during the resource-creating action. Therefore, a user who has permissions to create a resource (assuming there are no tagging conditions) does not require permission to use the `fsx:TagResource` action if no tags are specified in the request. However, if the user attempts to create a resource with tags, the request fails if the user does not have permissions to use the `fsx:TagResource` action.

For more information about tagging Amazon FSx resources, see [Tag your Amazon FSx resources](#). For more information about using tags to control access to Amazon FSx resources, see [Using tags to control access to your Amazon FSx resources](#).

Using tags to control access to your Amazon FSx resources

To control access to Amazon FSx resources and actions, you can use IAM policies based on tags. You can provide this control in two ways:

- You can control access to Amazon FSx resources based on the tags on those resources.
- You can control which tags can be passed in an IAM request condition.

For information about how to use tags to control access to Amazon resources, see [Controlling access using tags](#) in the *IAM User Guide*. For more information about tagging Amazon FSx resources

at creation, see [Grant permission to tag resources during creation](#). For more information about tagging resources, see [Tag your Amazon FSx resources](#).

Controlling access based on tags on a resource

To control which actions a user or role can perform on an Amazon FSx resource, you can use tags on the resource. For example, you might want to allow or deny specific API operations on a file system resource based on the key-value pair of the tag on the resource.

Example Example policy – Create a file system only when a specific tag is used

This policy allows the user to create a file system only when they tag it with a specific tag key-value pair, in this example, key=Department, value=Finance.

```
{
  "Effect": "Allow",
  "Action": [
    "fsx:CreateFileSystem",
    "fsx:TagResource"
  ],
  "Resource": "arn:aws:fsx:region:account-id:file-system/*",
  "Condition": {
    "StringEquals": {
      "aws:RequestTag/Department": "Finance"
    }
  }
}
```

Example Example policy – Create backups only of Amazon FSx for NetApp ONTAP volumes with a specific tag

This policy allows users to create backups only of FSx for ONTAP volumes that are tagged with the key-value pair key=Department, value=Finance. The backup is created with the tag Department=Finance.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
```



```

        "fsx:CreateBackup"
    ],
    "Resource": "arn:aws:fsx:region:account-id:volume/*",
    "Condition": {
        "StringEquals": {
            "aws:ResourceTag/Department": "Finance"
        }
    }
},
{
    "Effect": "Allow",
    "Action": [
        "fsx:TagResource",
        "fsx:CreateBackup"
    ],
    "Resource": "arn:aws:fsx:region:account-id:backup/*",
    "Condition": {
        "StringEquals": {
            "aws:RequestTag/Department": "Finance"
        }
    }
}
]
}

```

Example Example policy – Create a volume with a specific tag from backups with a specific tag

This policy allows users to create volumes that are tagged with Department=Finance only from backups that are tagged with Department=Finance.

```

{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "fsx:CreateVolumeFromBackup",
                "fsx:TagResource"
            ],
            "Resource": "arn:aws:fsx:region:account-id:volume/*",
            "Condition": {
                "StringEquals": {

```

```

        "aws:RequestTag/Department": "Finance"
      }
    }
  },
  {
    "Effect": "Allow",
    "Action": [
      "fsx:CreateVolumeFromBackup"
    ],
    "Resource": "arn:aws:fsx:region:account-id:backup/*",
    "Condition": {
      "StringEquals": {
        "aws:ResourceTag/Department": "Finance"
      }
    }
  }
]
}

```

Example Example policy – Delete file systems with specific tags

This policy allows a user to delete only file systems that are tagged with Department=Finance. If they create a final backup, then it must be tagged with Department=Finance.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "fsx>DeleteFileSystem"
      ],
      "Resource": "arn:aws:fsx:region:account-id:file-system/*",
      "Condition": {
        "StringEquals": {
          "aws:ResourceTag/Department": "Finance"
        }
      }
    },
    {
      "Effect": "Allow",
      "Action": [
        "fsx:TagResource"
      ]
    }
  ]
}

```

```

    ],
    "Resource": "arn:aws:fsx:region:account-id:backup/*",
    "Condition": {
      "StringEquals": {
        "aws:RequestTag/Department": "Finance"
      }
    }
  }
]
}

```

Example Example policy – Delete a volume with specific tags

This policy allows a user to delete only volumes that are tagged with Department=Finance. If they create a final backup, then it must be tagged with Department=Finance.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "fsx:DeleteVolume"
      ],
      "Resource": "arn:aws:fsx:region:account-id:volume/*",
      "Condition": {
        "StringEquals": {
          "aws:ResourceTag/Department": "Finance"
        }
      }
    },
    {
      "Effect": "Allow",
      "Action": [
        "fsx:TagResource"
      ],
      "Resource": "arn:aws:fsx:region:account-id:backup/*",
      "Condition": {
        "StringEquals": {
          "aws:RequestTag/Department": "Finance"
        }
      }
    }
  ]
}

```

```
]
}
```

Using service-linked roles for Amazon FSx

Amazon FSx uses Amazon Identity and Access Management (IAM) [service-linked roles](#). A service-linked role is a unique type of IAM role that is linked directly to Amazon FSx. Service-linked roles are predefined by Amazon FSx and include all the permissions that the service requires to call other Amazon services on your behalf.

A service-linked role makes setting up Amazon FSx easier because you don't have to manually add the necessary permissions. Amazon FSx defines the permissions of its service-linked roles, and unless defined otherwise, only Amazon FSx can assume its roles. The defined permissions include the trust policy and the permissions policy, and that permissions policy cannot be attached to any other IAM entity.

You can delete a service-linked role only after first deleting their related resources. This protects your Amazon FSx resources because you can't inadvertently remove permission to access the resources.

For information about other services that support service-linked roles, see [Amazon Services That Work with IAM](#) and look for the services that have **Yes** in the **Service-Linked Role** column. Choose a **Yes** with a link to view the service-linked role documentation for that service.

Service-linked role permissions for Amazon FSx

Amazon FSx uses the service-linked role named **AWSServiceRoleForAmazonFSx** – Which performs certain actions in your account, like creating Elastic Network Interfaces for your file systems in your VPC, and publishing file system and volume metrics in CloudWatch.

For updates to this policy, see [AmazonFSxServiceRolePolicy](#)

Permissions details

Permissions details

The **AWSServiceRoleForAmazonFSx** role permissions are defined by the **AmazonFSxServiceRolePolicy** Amazon managed policy. The **AWSServiceRoleForAmazonFSx** has the following permissions:

Note

The `AWSServiceRoleForAmazonFSx` is used by all Amazon FSx file system types; some of the listed permissions are not applicable to FSx for ONTAP.

- `ds` – Allows Amazon FSx to view, authorize, and unauthorize applications in your Amazon Directory Service directory.
- `ec2` – Allows Amazon FSx to do the following:
 - View, create, and disassociate network interfaces associated with an Amazon FSx file system.
 - View one or more Elastic IP addresses associated with an Amazon FSx file system.
 - View Amazon VPCs, security groups, and subnets associated with an Amazon FSx file system.
 - To provide enhanced security group validation of all security groups that can be used with a VPC.
 - Create a permission for an Amazon-authorized user to perform certain operations on a network interface.
- `ccloudwatch` – Allows Amazon FSx to publish metric data points to CloudWatch under the `Amazon/FSx` namespace.
- `route53` – Allows Amazon FSx to associate an Amazon VPC with a private hosted zone.
- `logs` – Allows Amazon FSx to describe and write to CloudWatch Logs log streams. This is so that users can send file access audit logs for an FSx for Windows File Server file system to a CloudWatch Logs stream.
- `firehose` – Allows Amazon FSx to describe and write to Amazon Data Firehose delivery streams. This is so that users can publish the file access audit logs for an Amazon FSx for Windows File Server file system to an Amazon Data Firehose delivery stream.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "CreateFileSystem",
      "Effect": "Allow",
      "Action": [
        "ds:AuthorizeApplication",
        "ds:GetAuthorizedApplicationDetails",
```

```

        "ds:UnauthorizeApplication",
        "ec2:CreateNetworkInterface",
        "ec2:CreateNetworkInterfacePermission",
        "ec2>DeleteNetworkInterface",
        "ec2:DescribeAddresses",
        "ec2:DescribeDhcpOptions",
        "ec2:DescribeNetworkInterfaces",
        "ec2:DescribeRouteTables",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeSubnets",
        "ec2:DescribeVPCs",
        "ec2:DisassociateAddress",
        "ec2:GetSecurityGroupsForVpc",
        "route53:AssociateVPCWithHostedZone"
    ],
    "Resource": "*"
},
{
    "Sid": "PutMetrics",
    "Effect": "Allow",
    "Action": [
        "cloudwatch:PutMetricData"
    ],
    "Resource": [
        "*"
    ],
    "Condition": {
        "StringEquals": {
            "cloudwatch:namespace": "AWS/FSx"
        }
    }
},
{
    "Sid": "TagResourceNetworkInterface",
    "Effect": "Allow",
    "Action": [
        "ec2:CreateTags"
    ],
    "Resource": [
        "arn:aws:ec2:*:*:network-interface/*"
    ],
    "Condition": {
        "StringEquals": {

```

```

        "ec2:CreateAction": "CreateNetworkInterface"
    },
    "ForAllValues:StringEquals": {
        "aws:TagKeys": "AmazonFSx.FileSystemId"
    }
}
},
{
    "Sid": "ManageNetworkInterface",
    "Effect": "Allow",
    "Action": [
        "ec2:AssignPrivateIpAddresses",
        "ec2:ModifyNetworkInterfaceAttribute",
        "ec2:UnassignPrivateIpAddresses"
    ],
    "Resource": [
        "arn:aws:ec2:*:*:network-interface/*"
    ],
    "Condition": {
        "Null": {
            "aws:ResourceTag/AmazonFSx.FileSystemId": "false"
        }
    }
},
{
    "Sid": "ManageRouteTable",
    "Effect": "Allow",
    "Action": [
        "ec2:CreateRoute",
        "ec2:ReplaceRoute",
        "ec2>DeleteRoute"
    ],
    "Resource": [
        "arn:aws:ec2:*:*:route-table/*"
    ],
    "Condition": {
        "StringEquals": {
            "aws:ResourceTag/AmazonFSx": "ManagedByAmazonFSx"
        }
    }
},
{
    "Sid": "PutCloudWatchLogs",
    "Effect": "Allow",

```

```

    "Action": [
      "logs:DescribeLogGroups",
      "logs:DescribeLogStreams",
      "logs:PutLogEvents"
    ],
    "Resource": "arn:aws:logs:*:*:log-group:/aws/fsx/*"
  },
  {
    "Sid": "ManageAuditLogs",
    "Effect": "Allow",
    "Action": [
      "firehose:DescribeDeliveryStream",
      "firehose:PutRecord",
      "firehose:PutRecordBatch"
    ],
    "Resource": "arn:aws:firehose:*:*:deliverystream/aws-fsx-*"
  }
]
}

```

Any updates to this policy are described in [Amazon FSx updates to Amazon managed policies](#).

You must configure permissions to allow an IAM entity (such as a user, group, or role) to create, edit, or delete a service-linked role. For more information, see [Service-Linked Role Permissions](#) in the IAM User Guide.

Creating a service-linked role for Amazon FSx

You don't need to manually create a service-linked role. When you create a file system in the Amazon Web Services Management Console, the IAM CLI, or the IAM API, Amazon FSx creates the service-linked role for you.

Important

This service-linked role can appear in your account if you completed an action in another service that uses the features supported by this role. To learn more, see [A New Role Appeared in My IAM Account](#).

If you delete this service-linked role, and then need to create it again, you can use the same process to recreate the role in your account. When you create a file system, Amazon FSx creates the service-linked role for you again.

Editing a service-linked role for Amazon FSx

Amazon FSx does not allow you to edit the `AWSServiceRoleForAmazonFSx` service-linked role. After you create a service-linked role, you cannot change the name of the role because various entities might reference the role. However, you can edit the description of the role using IAM. For more information, see [Editing a Service-Linked Role](#) in the *IAM User Guide*.

Deleting a service-linked role for Amazon FSx

If you no longer need to use a feature or service that requires a service-linked role, we recommend that you delete that role. That way you don't have an unused entity that is not actively monitored or maintained. However, you must delete all of your file systems and backups before you can manually delete the service-linked role.

Note

If the Amazon FSx service is using the role when you try to delete the resources, then the deletion might fail. If that happens, wait for a few minutes and try the operation again.

To manually delete the service-linked role using IAM

Use the IAM console, the IAM CLI, or the IAM API to delete the `AWSServiceRoleForAmazonFSx` service-linked role. For more information, see [Deleting a Service-Linked Role](#) in the *IAM User Guide*.

Supported regions for Amazon FSx service-linked roles

Amazon FSx supports using service-linked roles in all of the regions where the service is available. For more information, see [Amazon Regions and Endpoints](#).

Amazon managed policies for Amazon FSx

An Amazon managed policy is a standalone policy that is created and administered by Amazon. Amazon managed policies are designed to provide permissions for many common use cases so that you can start assigning permissions to users, groups, and roles.

Keep in mind that Amazon managed policies might not grant least-privilege permissions for your specific use cases because they're available for all Amazon customers to use. We recommend that you reduce permissions further by defining [customer managed policies](#) that are specific to your use cases.

You cannot change the permissions defined in Amazon managed policies. If Amazon updates the permissions defined in an Amazon managed policy, the update affects all principal identities (users, groups, and roles) that the policy is attached to. Amazon is most likely to update an Amazon managed policy when a new Amazon Web Service is launched or new API operations become available for existing services.

For more information, see [Amazon managed policies](#) in the *IAM User Guide*.

AmazonFSxServiceRolePolicy

Allows Amazon FSx to manage Amazon resources on your behalf. See [Using service-linked roles for Amazon FSx](#) to learn more.

Amazon managed policy: AmazonFSxDeleteServiceLinkedRoleAccess

You can't attach `AmazonFSxDeleteServiceLinkedRoleAccess` to your IAM entities. This policy is linked to a service and used only with the service-linked role for that service. You cannot attach, detach, modify, or delete this policy. For more information, see [Using service-linked roles for Amazon FSx](#).

This policy grants administrative permissions that allow Amazon FSx to delete its Service Linked Role for Amazon S3 access, used only by Amazon FSx for Lustre.

Permissions details

This policy includes permissions in `iam` to allow Amazon FSx to view, delete, and view the deletion status for the FSx Service Linked Roles for Amazon S3 access.

To view the permissions for this policy, see [AmazonFSxDeleteServiceLinkedRoleAccess](#) in the Amazon Managed Policy Reference Guide.

Amazon managed policy: AmazonFSxFullAccess

You can attach `AmazonFSxFullAccess` to your IAM entities. Amazon FSx also attaches this policy to a service role that allows Amazon FSx to perform actions on your behalf.

Provides full access to Amazon FSx and access to related Amazon services.

Permissions details

This policy includes the following permissions.

- `fsx` – Allows principals full access to perform all Amazon FSx actions, except for `BypassSnaplockEnterpriseRetention`.
- `ds` – Allows principals to view information about the Amazon Directory Service directories.
- `ec2`
 - Allows principals to create tags under the specified conditions.
 - To provide enhanced security group validation of all security groups that can be used with a VPC.
- `iam` – Allows principles to create an Amazon FSx service linked role on the user's behalf. This is required so that Amazon FSx can manage Amazon resources on the user's behalf.
- `logs` – Allows principals to create log groups, log streams, and write events to log streams. This is required so that users can monitor FSx for Windows File Server file system access by sending audit access logs to CloudWatch Logs.
- `firehose` – Allows principals to write records to a Amazon Data Firehose. This is required so that users can monitor FSx for Windows File Server file system access by sending audit access logs to Firehose.

To view the permissions for this policy, see [AmazonFSxFullAccess](#) in the Amazon Managed Policy Reference Guide.

Amazon managed policy: AmazonFSxConsoleFullAccess

You can attach the `AmazonFSxConsoleFullAccess` policy to your IAM identities.

This policy grants administrative permissions that allow full access to Amazon FSx and access to related Amazon services via the Amazon Web Services Management Console.

Permissions details

This policy includes the following permissions.

- `fsx` – Allows principals to perform all actions in the Amazon FSx management console, except for `BypassSnaplockEnterpriseRetention`.
- `cloudwatch` – Allows principals to view CloudWatch Alarms and metrics in the Amazon FSx management console.
- `ds` – Allows principals to list information about an Amazon Directory Service directory.
- `ec2`
 - Allows principals to create tags on route tables, list network interfaces, route tables, security groups, subnets and the VPC associated with an Amazon FSx file system.
 - Allows principals to provide enhanced security group validation of all security groups that can be used with a VPC.
- `kms` – Allows principals to list aliases for Amazon Key Management Service keys.
- `s3` – Allows principals to list some or all of the objects in an Amazon S3 bucket (up to 1000).
- `iam` – Grants permission to create a service linked role that allows Amazon FSx to perform actions on the user's behalf.

To view the permissions for this policy, see [AmazonFSxConsoleFullAccess](#) in the Amazon Managed Policy Reference Guide.

Amazon managed policy: AmazonFSxConsoleReadOnlyAccess

You can attach the `AmazonFSxConsoleReadOnlyAccess` policy to your IAM identities.

This policy grants read-only permissions to Amazon FSx and related Amazon services so that users can view information about these services in the Amazon Web Services Management Console.

Permissions details

This policy includes the following permissions.

- `fsx` – Allows principals to view information about Amazon FSx file systems, including all tags, in the Amazon FSx Management Console.
- `cloudwatch` – Allows principals to view CloudWatch Alarms and metrics in the Amazon FSx Management Console.
- `ds` – Allows principals to view information about an Amazon Directory Service directory in the Amazon FSx Management Console.

- `ec2`
 - Allows principals to view network interfaces, security groups, subnets and the VPC associated with an Amazon FSx file system in the Amazon FSx Management Console.
 - To provide enhanced security group validation of all security groups that can be used with a VPC.
- `kms` – Allows principals to view aliases for Amazon Key Management Service keys in the Amazon FSx Management Console.
- `log` – Allows principals to describe the Amazon CloudWatch Logs log groups associated with the account making the request. This is required so that principals can view the existing file access auditing configuration for an FSx for Windows File Server file system.
- `firehose` – Allows principals to describe the Amazon Data Firehose delivery streams associated with the account making the request. This is required so that principals can view the existing file access auditing configuration for an FSx for Windows File Server file system.

To view the permissions for this policy, see [AmazonFSxConsoleReadOnlyAccess](#) in the Amazon Managed Policy Reference Guide.

Amazon managed policy: AmazonFSxReadOnlyAccess

You can attach the AmazonFSxReadOnlyAccess policy to your IAM identities.

This policy includes the following permissions.

- `fsx` – Allows principals to view information about Amazon FSx file systems, including all tags, in the Amazon FSx Management Console.
- `ec2` – To provide enhanced security group validation of all security groups that can be used with a VPC.

To view the permissions for this policy, see [AmazonFSxReadOnlyAccess](#) in the Amazon Managed Policy Reference Guide.

Amazon FSx updates to Amazon managed policies

View details about updates to Amazon managed policies for Amazon FSx since this service began tracking these changes. For automatic alerts about changes to this page, subscribe to the RSS feed on the Amazon FSx [Document History for Amazon FSx for NetApp ONTAP](#) page.

Change	Description	Date
AmazonFSxServiceRolePolicy – Update to an existing policy	Amazon FSx added new permission, <code>ec2:GetSecurityGroupsForVpc</code> that allows principals to provide enhanced security group validation of all security groups that can be used with a VPC.	January 9, 2024
AmazonFSxReadOnlyAccess – Update to an existing policy	Amazon FSx added new permission, <code>ec2:GetSecurityGroupsForVpc</code> that allows principals to provide enhanced security group validation of all security groups that can be used with a VPC.	January 9, 2024
AmazonFSxConsoleReadOnlyAccess – Update to an existing policy	Amazon FSx added new permission, <code>ec2:GetSecurityGroupsForVpc</code> that allows principals to provide enhanced security group validation of all security groups that can be used with a VPC.	January 9, 2024
AmazonFSxFullAccess – Update to an existing policy	Amazon FSx added new permission, <code>ec2:GetSecurityGroupsForVpc</code> that allows principals to provide enhanced security group validation of all security groups that can be used with a VPC.	January 9, 2024

Change	Description	Date
AmazonFSxConsoleFullAccess – Update to an existing policy	Amazon FSx added new permission, <code>ec2:GetSecurityGroupsForVpc</code> that allows principals to provide enhanced security group validation of all security groups that can be used with a VPC.	January 9, 2024
AmazonFSxFullAccess – Update to an existing policy	Amazon FSx added new permission to enable users to perform cross-region and cross-account data replication for FSx for OpenZFS file systems.	December 20, 2023
AmazonFSxConsoleFullAccess – Update to an existing policy	Amazon FSx added new permission to enable users to perform cross-region and cross-account data replication for FSx for OpenZFS file systems.	December 20, 2023
AmazonFSxFullAccess – Update to an existing policy	Amazon FSx added a new permission to enable users to perform on-demand replication of volumes for FSx for OpenZFS file systems.	November 26, 2023
AmazonFSxConsoleFullAccess – Update to an existing policy	Amazon FSx added a new permission to enable users to perform on-demand replication of volumes for FSx for OpenZFS file systems.	November 26, 2023

Change	Description	Date
AmazonFSxFullAccess – Update to an existing policy	Amazon FSx added new permissions to enable users to view, enable, and disable shared VPC support for FSx for ONTAP Multi-AZ file systems.	November 14, 2023
AmazonFSxConsoleFullAccess – Update to an existing policy	Amazon FSx added new permissions to enable users to view, enable, and disable shared VPC support for FSx for ONTAP Multi-AZ file systems.	November 14, 2023
AmazonFSxFullAccess – Update to an existing policy	Amazon FSx added new permissions to allow Amazon FSx to manage network configurations for FSx for OpenZFS Multi-AZ file systems.	August 9, 2023
Amazon managed policy: AmazonFSxServiceRolePolicy – Update to an existing policy	Amazon FSx modified the existing <code>cloudwatch:PutMetricData</code> permission so that Amazon FSx publishes CloudWatch metrics to the Amazon/FSx namespace.	July 24, 2023
AmazonFSxFullAccess – Update to an existing policy	Amazon FSx updated the policy to remove the <code>fsx:*</code> permission and add specific fsx actions.	July 13, 2023

Change	Description	Date
AmazonFSxConsoleFullAccess – Update to an existing policy	Amazon FSx updated the policy to remove the <code>fsx:*</code> permission and add specific <code>fsx</code> actions.	July 13, 2023
AmazonFSxConsoleReadOnlyAccess – Update to an existing policy	Amazon FSx added new permissions to enable users to view enhanced performance metrics and recommended actions for FSx for Windows File Server file systems in the Amazon FSx console.	September 21, 2022
AmazonFSxConsoleFullAccess – Update to an existing policy	Amazon FSx added new permissions to enable users to view enhanced performance metrics and recommended actions for FSx for Windows File Server file systems in the Amazon FSx console.	September 21, 2022
AmazonFSxReadOnlyAccess – Started tracking policy	This policy grants read-only access to all Amazon FSx resources and any tags associated with them.	February 4, 2022
AmazonFSxDeleteServiceLinkedRoleAccess – Started tracking policy	This policy grants administrative permissions that allow Amazon FSx to delete its Service Linked Role for Amazon S3 access.	January 7, 2022

Change	Description	Date
AmazonFSxServiceRolePolicy – Update to an existing policy	Amazon FSx added new permissions to allow Amazon FSx to manage network configurations for Amazon FSx for NetApp ONTAP file systems.	September 2, 2021
AmazonFSxFullAccess – Update to an existing policy	Amazon FSx added new permissions to allow Amazon FSx to create tags on EC2 route tables for scoped down calls.	September 2, 2021
AmazonFSxConsoleFullAccess – Update to an existing policy	Amazon FSx added new permissions to allow Amazon FSx to create Amazon FSx for NetApp ONTAP Multi-AZ file systems.	September 2, 2021
AmazonFSxConsoleFullAccess – Update to an existing policy	Amazon FSx added new permissions to allow Amazon FSx to create tags on EC2 route tables for scoped down calls.	September 2, 2021
AmazonFSxServiceRolePolicy – Update to an existing policy	Amazon FSx added new permissions to allow Amazon FSx to describe and write to CloudWatch Logs log streams. This is required so that users can view file access audit logs for FSx for Windows File Server file systems using CloudWatch Logs.	June 8, 2021

Change	Description	Date
AmazonFSxServiceRolePolicy – Update to an existing policy	<p>Amazon FSx added new permissions to allow Amazon FSx to describe and write to Amazon Data Firehose delivery streams.</p> <p>This is required so that users can view file access audit logs for an FSx for Windows File Server file system using Amazon Data Firehose.</p>	June 8, 2021
AmazonFSxFullAccess – Update to an existing policy	<p>Amazon FSx added new permissions to allow principals to describe and create CloudWatch Logs log groups, log streams, and write events to log streams.</p> <p>This is required so that principals can view file access audit logs for FSx for Windows File Server file systems using CloudWatch Logs.</p>	June 8, 2021

Change	Description	Date
AmazonFSxFullAccess – Update to an existing policy	<p>Amazon FSx added new permissions to allow principals to describe and write records to a Amazon Data Firehose.</p> <p>This is required so that users can view file access audit logs for an FSx for Windows File Server file system using Amazon Data Firehose.</p>	June 8, 2021
AmazonFSxConsoleFullAccess – Update to an existing policy	<p>Amazon FSx added new permissions to allow principals to describe the Amazon CloudWatch Logs log groups associated with the account making the request.</p> <p>This is required so that principals can choose an existing CloudWatch Logs log group when configuring file access auditing for an FSx for Windows File Server file system.</p>	June 8, 2021

Change	Description	Date
<p>AmazonFSxConsoleFullAccess – Update to an existing policy</p>	<p>Amazon FSx added new permissions to allow principals to describe the Amazon Data Firehose delivery streams associated with the account making the request.</p> <p>This is required so that principals can choose an existing Firehose delivery stream when configuring file access auditing for an FSx for Windows File Server file system.</p>	<p>June 8, 2021</p>
<p>AmazonFSxConsoleReadOnlyAccess – Update to an existing policy</p>	<p>Amazon FSx added new permissions to allow principals to describe the Amazon CloudWatch Logs log groups associated with the account making the request.</p> <p>This is required so that principals can view the existing file access auditing configuration for an FSx for Windows File Server file system.</p>	<p>June 8, 2021</p>

Change	Description	Date
AmazonFSxConsoleReadOnlyAccess – Update to an existing policy	<p>Amazon FSx added new permissions to allow principals to describe the Amazon Data Firehose delivery streams associated with the account making the request.</p> <p>This is required so that principals can view the existing file access auditing configuration for an FSx for Windows File Server file system.</p>	June 8, 2021
Amazon FSx started tracking changes	Amazon FSx started tracking changes for its Amazon managed policies.	June 8, 2021

File System Access Control with Amazon VPC

You access your Amazon FSx for NetApp ONTAP file systems and SVMs using the DNS name or the IP address of one of their endpoints, depending on what type of access it is. The DNS name maps to the private IP address of the file system's or SVM's elastic network interface in your VPC. Only resources within the associated VPC, or resources connected with the associated VPC by Amazon Direct Connect or VPN, can access the data in your file system over the NFS, SMB, or iSCSI protocols. For more information, see [What is Amazon VPC?](#) in the *Amazon VPC User Guide*.

Warning

You must not modify or delete the elastic network interface(s) associated with your file system. Modifying or deleting the network interface can cause a permanent loss of connection between your VPC and your file system.

Amazon VPC security groups

A security group acts as a virtual firewall for your FSx for ONTAP file systems to control incoming and outgoing traffic. Inbound rules control the incoming traffic to your file system, and outbound rules control the outgoing traffic from your file system. When you create a file system, you specify the VPC that it gets created in, and the default security group for that VPC is applied. You can add rules to each security group that allow traffic to or from its associated file systems and SVMs. You can modify the rules for a security group at any time. New and modified rules are automatically applied to all resources that are associated with the security group. When Amazon FSx decides whether to allow traffic to reach a resource, it evaluates all of the rules from all of the security groups that are associated with the resource.

To use a security group to control access to your Amazon FSx file system, add inbound and outbound rules. Inbound rules control incoming traffic, and outbound rules control outgoing traffic from your file system. Make sure that you have the right network traffic rules in your security group to map your Amazon FSx file system's file share to a folder on your supported compute instance.

For more information on security group rules, see [Security Group Rules](#) in the *Amazon EC2 User Guide for Linux Instances*.

Creating a VPC security group

To create a security group for Amazon FSx

1. Open the Amazon EC2 console at <https://console.amazonaws.cn/ec2>.
2. In the navigation pane, choose **Security Groups**.
3. Choose **Create Security Group**.
4. Specify a name and description for the security group.
5. For **VPC**, choose the Amazon VPC associated with your file system to create the security group within that VPC.
6. For outbound rules, allow all traffic on all ports.
7. Add the following rules to the inbound ports of your security group. For the **source** field, you should choose **Custom** and enter the security groups or IP address ranges associated with the instances that need to access your FSx for ONTAP file system, including:
 - Linux, Windows, and/or macOS clients that access data in your file system over NFS, SMB, or iSCSI.

- Any ONTAP file systems/clusters that you will peer to your file system (for example, to use SnapMirror, SnapVault, or FlexCache).
- Any clients that you will use to access the ONTAP REST API, CLI, or ZAPIs (for example, a Harvest/Grafana instance, NetApp Connector, or NetApp BlueXP).

Protocol	Ports	Role
All ICMP	All	Pinging the instance
SSH	22	SSH access to the IP address of the cluster management LIF or a node management LIF
TCP	111	Remote procedure call for NFS
TCP	135	Remote procedure call for CIFS
TCP	139	NetBIOS service session for CIFS
TCP	161-162	Simple network management protocol (SNMP)
TCP	443	ONTAP REST API access to the IP address of the cluster management LIF or an SVM management LIF
TCP	445	Microsoft SMB/CIFS over TCP with NetBIOS framing
TCP	635	NFS mount
TCP	749	Kerberos
TCP	2049	NFS server daemon
TCP	3260	iSCSI access through the iSCSI data LIF
TCP	4045	NFS lock daemon
TCP	4046	Network status monitor for NFS
TCP	10000	Network data management protocol (NDMP) and NetApp SnapMirror intercluster communication

Protocol	Ports	Role
TCP	11104	Management of NetApp SnapMirror intercluster communication
TCP	11105	SnapMirror data transfer using intercluster LIFs
UDP	111	Remote procedure call for NFS
UDP	135	Remote procedure call for CIFS
UDP	137	NetBIOS name resolution for CIFS
UDP	139	NetBIOS service session for CIFS
UDP	161-162	Simple network management protocol (SNMP)
UDP	635	NFS mount
UDP	2049	NFS server daemon
UDP	4045	NFS lock daemon
UDP	4046	Network status monitor for NFS
UDP	4049	NFS quota protocol

8. Add the security group to the file system's elastic network interface.

Disallow access to a file system

To temporarily disallow network access to your file system from all clients, you can remove all the security groups associated with your file system's elastic network interface(s) and replace them with a group that has no inbound/outbound rules.

Compliance Validation for Amazon FSx for NetApp ONTAP

To learn whether an Amazon Web Service is within the scope of specific compliance programs, see [Amazon Web Services in Scope by Compliance Program](#) and choose the compliance program that you are interested in. For general information, see [Amazon Web Services Compliance Programs](#).

You can download third-party audit reports using Amazon Artifact. For more information, see [Downloading Reports in Amazon Artifact](#).

Your compliance responsibility when using Amazon Web Services is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. Amazon provides the following resources to help with compliance:

- [Security and Compliance Quick Start Guides](#) – These deployment guides discuss architectural considerations and provide steps for deploying baseline environments on Amazon that are security and compliance focused.
- [Amazon Compliance Resources](#) – This collection of workbooks and guides might apply to your industry and location.
- [Evaluating Resources with Rules](#) in the *Amazon Config Developer Guide* – The Amazon Config service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- [Amazon Security Hub](#) – This Amazon Web Service provides a comprehensive view of your security state within Amazon. Security Hub uses security controls to evaluate your Amazon resources and to check your compliance against security industry standards and best practices. For a list of supported services and controls, see [Security Hub controls reference](#).

Amazon FSx for NetApp ONTAP and interface VPC endpoints (Amazon PrivateLink)

You can improve the security posture of your VPC by configuring Amazon FSx to use an interface VPC endpoint. Interface VPC endpoints are powered by [Amazon PrivateLink](#), a technology that enables you to privately access Amazon FSx APIs without an internet gateway, NAT device, VPN connection, or Amazon Direct Connect connection. Instances in your VPC don't need public IP addresses to communicate with Amazon FSx APIs. Traffic between your VPC and Amazon FSx does not leave the Amazon network.

Each interface VPC endpoint is represented by one or more elastic network interfaces in your subnets. A network interface provides a private IP address that serves as an entry point for traffic to the Amazon FSx API.

Considerations for Amazon FSx interface VPC endpoints

Before you set up an interface VPC endpoint for Amazon FSx, be sure to review [Interface VPC endpoint properties and limitations](#) in the *Amazon VPC User Guide*.

You can call any of the Amazon FSx API operations from your VPC. For example, you can create an FSx for ONTAP file system by calling the CreateFileSystem API from within your VPC. For the full list of Amazon FSx APIs, see [Actions](#) in the Amazon FSx API Reference.

VPC peering considerations

You can connect other VPCs to the VPC with interface VPC endpoints using VPC peering. VPC peering is a networking connection between two VPCs. You can establish a VPC peering connection between your own two VPCs, or with a VPC in another Amazon Web Services account. The VPCs can also be in two different Amazon Web Services Regions.

Traffic between peered VPCs stays on the Amazon network and does not traverse the public internet. Once VPCs are peered, resources like Amazon Elastic Compute Cloud (Amazon EC2) instances in both VPCs can access the Amazon FSx API through interface VPC endpoints created in the one of the VPCs.

Creating an interface VPC endpoint for Amazon FSx API

You can create a VPC endpoint for the Amazon FSx API using either the Amazon VPC console or the Amazon Command Line Interface (Amazon CLI). For more information, see [Creating an interface VPC endpoint](#) in the *Amazon VPC User Guide*.

To create an interface VPC endpoint for Amazon FSx, use one of the following:

- **com.amazonaws.*region*.fsx** – Creates an endpoint for Amazon FSx API operations.
- **com.amazonaws.*region*.fsx-fips** – Creates an endpoint for the Amazon FSx API that complies with [Federal Information Processing Standard \(FIPS\) 140-2](#).

To use the private DNS option, you must set the `enableDnsHostnames` and `enableDnsSupport` attributes of your VPC. For more information, see [Viewing and updating DNS support for your VPC](#) in the *Amazon VPC User Guide*.

Excluding Amazon Web Services Regions in China, if you enable private DNS for the endpoint, you can make API requests to Amazon FSx with the VPC endpoint using its default DNS name

for the Amazon Web Services Region, for example `fsx.us-east-1.amazonaws.com`. For the China (Beijing) and China (Ningxia) Amazon Web Services Regions, you can make API requests with the VPC endpoint using `fsx-api.cn-north-1.amazonaws.com.cn` and `fsx-api.cn-northwest-1.amazonaws.com.cn`, respectively.

For more information, see [Accessing a service through an interface VPC endpoint](#) in the *Amazon VPC User Guide*.

Creating a VPC endpoint policy for Amazon FSx

To control access to the Amazon FSx API, you can attach an Amazon Identity and Access Management (IAM) policy to your VPC endpoint. The policy specifies the following:

- The principal that can perform actions.
- The actions that can be performed.
- The resources on which actions can be performed.

For more information, see [Controlling access to services with VPC endpoints](#) in the *Amazon VPC User Guide*.

Resilience in Amazon FSx for NetApp ONTAP

The Amazon global infrastructure is built around Amazon Web Services Regions and Availability Zones. Amazon Web Services Regions provide multiple physically separated and isolated Availability Zones, which are connected with low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For more information about Amazon Web Services Regions and Availability Zones, see [Amazon Global Infrastructure](#).

In addition to the Amazon global infrastructure, Amazon FSx offers several features to help support your data resiliency and backup needs.

Backup and restore

Amazon FSx creates and saves automated backups of the volumes in your Amazon FSx for NetApp ONTAP file system. Amazon FSx creates automated backups of your volumes during the backup

window of your Amazon FSx for NetApp ONTAP file system. Amazon FSx saves the automated backups of your volumes according to the backup retention period that you specify. You can also back up your volumes manually, by creating a user-initiated backup. You restore a volume backup at any time by creating a new volume with the backup specified as the source.

For more information, see [Working with backups](#).

Snapshots

Amazon FSx creates snapshot copies of the Amazon FSx for NetApp ONTAP volumes. Snapshot copies offer protection against accidental deletion or modification of files in your volumes by end users. For more information, see [Working with snapshots](#).

Availability Zones

Amazon FSx for NetApp ONTAP file systems are designed to provide continuous availability to data even in the event that a server failure. Each file system is powered by two file servers in at least one Availability Zone, each with its own storage. Amazon FSx automatically replicates your data to protect it from component failure, continuously monitors for hardware failures, and automatically replaces infrastructure components in the event of a failure. File systems automatically fail over and back as needed (typically within 60 seconds), and clients automatically fail over and back with the file system.

Multi-AZ file systems

Amazon FSx for NetApp ONTAP file systems are highly available and durable across Amazon Availability Zones, and are designed to provide continuous availability to data even in the event that an Availability Zone is unavailable.

For more information, see [Availability and durability](#).

Single-AZ file systems

Amazon FSx for NetApp ONTAP file systems are highly available and durable within a single Amazon Availability Zone, and are designed to provide continuous availability within that Availability Zone in the event of an individual file server or disk failure.

For more information, see [Availability and durability](#).

Infrastructure security in Amazon FSx for NetApp ONTAP

As a managed service, Amazon FSx for NetApp ONTAP is protected by Amazon global network security. For information about Amazon security services and how Amazon protects infrastructure, see [Amazon Cloud Security](#). To design your Amazon environment using the best practices for infrastructure security, see [Infrastructure Protection](#) in *Security Pillar Amazon Well-Architected Framework*.

You use Amazon published API calls to access Amazon FSx through the network. Clients must support the following:

- Transport Layer Security (TLS). We require TLS 1.2 and recommend TLS 1.3.
- Cipher suites with perfect forward secrecy (PFS) such as DHE (Ephemeral Diffie-Hellman) or ECDHE (Elliptic Curve Ephemeral Diffie-Hellman). Most modern systems such as Java 7 and later support these modes.

Additionally, requests must be signed by using an access key ID and a secret access key that is associated with an IAM principal. Or you can use the [Amazon Security Token Service](#) (Amazon STS) to generate temporary security credentials to sign requests.

Use NetApp ONTAP Vscan with FSx for ONTAP

You can use NetApp ONTAP's Vscan feature to run supported third-party antivirus software. For more information, see the following resources for each of the supported solutions.

- McAfee – [Antivirus Solution Guide for Clustered Data ONTAP: McAfee](#)
- SentinelOne – [Vscan partner solutions](#) and [SentinelOne Singularity Cloud Data Security](#)
- Symantec – [Vscan partner solutions](#) and [Symantec Protection Engine](#)
- Trend Micro – [Antivirus Solution Guide for Clustered Data ONTAP: Trend Micro](#)

Roles and users in Amazon FSx for NetApp ONTAP

ONTAP includes a robust and extensible role-based access control (RBAC) capability. You can assign users a role to control their access to the resources exposed through the ONTAP REST API and CLI. The roles define different levels of administrative access for the various ONTAP users. You can use roles and users in FSx for ONTAP to define user capabilities and privileges when using the ONTAP

CLI and REST API. Every ONTAP role and user is associated with your file system or a Storage Virtual Machine (SVM).

By default, your FSx for ONTAP file system has one file system-level user called *fsxadmin*, which has the `fsxadmin` role assigned to it. At the file system-level, you can only create new users with the `fsxadmin` role. You can't create new roles or modify the `fsxadmin` role.

For each SVM on your file system, there is a default user called *vsadmin*, which has the `vsadmin` role assigned. At the SVM-level, you can create new users and new roles. In addition to the `vsadmin` role, there are several predefined SVM roles you can assign to SVM users. You can also create roles that provide the level of access control that meet your organization's needs.

Predefined roles on an SVM

SVMs have the following predefined roles:

Role name	Capabilities
<code>vsadmin</code>	<ul style="list-style-type: none">• Manage your user account, local password, and key information• Manage volumes, except for volume moves• Manage quotas, qtrees, Snapshot copies, and files• Manage LUNs• Perform SnapLock operations, except for privileged delete• Configure protocols: NFS, SMB, and iSCSI• Configure services: DNS, LDAP, and NIS• Monitor jobs• Monitor network connections and the network interface• Monitor the health of the SVM
<code>vsadmin-volume</code>	<ul style="list-style-type: none">• Manage your user account, local password, and key information• Manage volumes, including volume moves

Role name	Capabilities
	<ul style="list-style-type: none"> • Manage quotas, qtrees, Snapshot copies, and files • Manage LUNs • Configure protocols: NFS, SMB, and iSCSI • Configure services: DNS, LDAP, and NIS • Monitor the network interface • Monitor the health of the SVM
vsadmin-protocol	<ul style="list-style-type: none"> • Manage your user account, local password, and key information • Manage LUNs • Configure protocols: NFS, SMB, and iSCSI • Configure services: DNS, LDAP, and NIS • Monitor network interface • Monitor the health of the SVM
vsadmin-backup	<ul style="list-style-type: none"> • Manage your user account, local password, and key information • Manage NDMP operations • Make a restored volume read/write • Manage SnapMirror relationships and Snapshot copies • View volumes and network information

Role name	Capabilities
vsadmin-snaplock	<ul style="list-style-type: none">• Manage your user account, local password, and key information• Manage volumes, except for volume moves• Manage quotas, qtrees, Snapshot copies, and files• Perform SnapLock operations, including privileged delete• Configure protocols: NFS and SMB• Configure services: DNS, LDAP, and NIS• Monitor jobs• Monitor network connections and the network interface
vsadmin-readonly	<ul style="list-style-type: none">• Manage your user account, local password, and key information• Monitor the health of the SVM• Monitor the network interface• View volumes and LUNs• View services and protocols

Topics

- [Creating new roles or users](#)
- [Updating the fsxadmin account password fails](#)
- [Creating new roles for an SVM using the NetApp ONTAP CLI](#)
- [Using Active Directory user accounts with your file system](#)
- [Configuring public key authentication](#)

Creating new roles or users

Each user in FSx for ONTAP is associated with an SVM or with the file system itself. You can create new roles or users by using the `security login create` command in the NetApp ONTAP CLI with the `vsadmin` role for SVMs and the default `fsxadmin` role for your file system.

The `security login create` command creates a login method for the management utility. A login method consists of a user name, an application (access method), and an authentication method. A user name can be associated with multiple applications. It can optionally include an access-control role name. If an Active Directory, LDAP, or NIS group name is used, then the login method gives access to users belonging to the specified group. If the user is a member of multiple groups provisioned in the security login table, then the user will get access to a combined list of the commands authorized for the individual groups. For more information, see the [security login create](#) in the NetApp ONTAP Product Documentation.

To create a new user for an SVM or file system (NetApp ONTAP CLI)

1. To access the NetApp ONTAP CLI, establish an SSH session on the management port of the Amazon FSx for NetApp ONTAP file system by running the following command. Replace *management_endpoint_ip* with the IP address of the file system's management port.

```
[~]$ ssh fsxadmin@management_endpoint_ip
```

For more information, see [Managing file systems with the ONTAP CLI](#).

2. Use the `security login create` ONTAP CLI command to create a new user account on your FSx for ONTAP file system or SVM.

```
Fsx0123456::> security login create -vserver vserver_name -user-or-group-name user_or_group_name -application login_application -authentication-method auth_method -role role_or_account_name
```

Insert your data for the placeholders in the example to define the following required properties:

- `-vserver` – Specifies the name of the file system or SVM where you want to create the new role or user.

- `-user-or-group-name` – Specifies the username or Active Directory group name of the login method. The Active Directory group name can be specified only with the `domain` authentication method and the `ontapi` and `ssh` applications.
- `-application` – Specifies the application of the login method. Possible values include `http`, `ontapi`, and `ssh`.
- `-authentication-method` – Specifies the authentication method for login. Possible values include the following:
 - `domain` – Use for Active Directory authentication
 - `password` – Use for password authentication
 - `publickey` – User for public-key authentication
- `-role` – Specifies the access-control role name for the login method. At the file system-level, the only role that can be specified is `fsxadmin`.

(Optional) You can also use one or more of the following parameters with the command:

- `[-comment]` – The comment text for the user account. For example, **Guest account**. The maximum length is 128 characters.
- `[-second-authentication-method {none|publickey|password|nsswitch}]` – Specifies the second factor authentication method. You can specify the following methods:
 - `password` – Use for password authentication
 - `publickey` – Use for Public-key authentication
 - `nsswitch` – Use for NIS or LDAP authentication
 - `none` – The default value if you don't specify one

Updating the `fsxadmin` account password fails

When you update the password for the `fsxadmin` user, you may receive an error if it doesn't meet the password requirements set on the file system. You can view the password requirements by using the `security login role config show` ONTAP CLI or REST API command.

To view the password requirements for the fsxadmin account

1. To access the NetApp ONTAP CLI, establish an SSH session on the management port of the Amazon FSx for NetApp ONTAP file system by running the following command. Replace *management_endpoint_ip* with the IP address of the file system's management port.

```
[~]$ ssh fsxadmin@management_endpoint_ip
```

For more information, see [Managing file systems with the ONTAP CLI](#).

2. The security login role config show command returns the password requirements for an account.

```
FsxId0123456::> security login role config show -role fsxadmin -  
fields password_requirement_fields
```

For the `-fields` parameter, specify any or all of the following:

- `passwd-minlength` – The minimum length of the password.
 - `passwd-min-special-chars` – The minimum number of special characters in the password.
 - `passwd-min-lowercase-chars` – The minimum number of lowercase characters in the password.
 - `passwd-min-uppercase-chars` – The minimum number of uppercase characters in the password.
 - `passwd-min-digits` – The minimum number of digits in the password.
 - `passwd-alphanum` – Information about the inclusion or exclusion of alphanumeric characters.
 - `passwd-expiry-time` – The password expiration time.
 - `passwd-expiry-warn-time` – The password expiration warning time.
3. Run the following command to see all password requirements:

```
Fsx0123456::> security login role config show -role fsxadmin -fields passwd-  
minlength, passwd-min-special-chars, passwd-min-lowercase-chars, passwd-min-  
digits, passwd-alphanum, passwd-expiry-time, passwd-expiry-warn-time, passwd-min-  
uppercase-chars
```

You will get a response similar to the following example, which displays the information for the specified fields.

```

vserver                role      passwd-minlength passwd-alphanum passwd-min-
special-chars passwd-expiry-time passwd-min-lowercase-chars passwd-min-uppercase-
chars passwd-min-digits passwd-expiry-warn-time
-----
-----
-----
FsxId0ae30e5b7f1a50b6a fsxadmin 3          enabled      0
          unlimited      0          0          0
          unlimited

```

To modify password requirements (ONTAP CLI)

- To modify your password requirements, use the `security login role config modify` command with the desired password fields. The following example shows how to change the minimum length of the password to 4 and the minimum number of special characters to 1 for the `fsxadmin` role on a file system.

```

Fsx0123456::> security login role config modify -role fsxadmin -passwd-minlength 4
               -passwd-min-special-chars 1

```

Creating new roles for an SVM using the NetApp ONTAP CLI

Each SVM that you create has a default SVM administrator that's assigned the predefined `vsadmin` role, but you can't create new roles using `vsadmin`. If you need to create new roles for your SVM, use the `security login role create` command with the `fsxadmin` role in the NetApp ONTAP CLI.

To create a new role on your SVM using the NetApp ONTAP CLI

- Copy the following `security login role create` command:

```

Fsx0123456::> security login role create -role vol_role -cmddirname "volume"

```

- Specify the following required parameters in the command:

- `-role` – The name of the role.
 - `-cmddirname` – The command or command directory to which the role gives access. Enclose command subdirectory names in double quotation marks. For example, "volume snapshot". Enter `DEFAULT` to specify all command directories.
3. (Optional) You can also add any of the following parameters to the command:
- `-vserver` – The name of the SVM that's associated with the role.
 - `-access` – The access level for the role. For command directories, this includes:
 - `none` – Denies access to commands in the command directory. This is the default value for custom roles.
 - `readonly` – Grants access to the show commands in the command directory and its subdirectories.
 - `all` – Grants access to all of the commands in the command directory and its subdirectories. To grant or deny access to intrinsic commands, you must specify the command directory.

For non-intrinsic commands (commands that don't end in `create`, `modify`, `delete`, or `show`):

- `none` – Denies access to commands in the command directory. This is the default value for custom roles.
 - `readonly` – Not applicable. Don't use.
 - `all` – Grants access to the command.
 - `-query` – The query object that's used to filter the access level, which is specified in the form of a valid option for the command, or for a command in the command directory. Enclose the query object in double quotation marks.
4. Run the `security login role create` command.

Using Active Directory user accounts with your file system

If you're a file system administrator, you can use the `security login domain-tunnel create` command in the NetApp ONTAP CLI to authenticate to your file system and SVMs using Active Directory accounts.

To authenticate to your file system or SVMs using Active Directory

1. To SSH into the NetApp ONTAP CLI of your file system, follow the steps documented in the [Using the NetApp ONTAP CLI](#) section of the FSx for ONTAP user guide.

```
ssh fsxadmin@file-system-management-endpoint-ip-address
```

2. To authenticate to your file system using Active Directory credentials instead of an individual SVM, copy the following `security login domain-tunnel create` command to set up domain tunneling. Select an SVM that's joined to your Active Directory to serve as the domain tunnel for authenticating logins with your Active Directory.

```
FsxIdabcdef0123456789a:> security login domain-tunnel create -vserver svm01
```

3. Use the `security login create` command to create one or more Active Directory domain user accounts that will be granted access to the file system.
4. Specify the following required parameters in the command:
 - `-vserver` – The name of the file system or SVM in which the new role or user will be created.
 - `-user-or-group-name` – The username or Active Directory group name of the login method. The Active Directory group name can be specified only with the `domain` authentication method and `ontapi` and `ssh` application.
 - `-application` – The application of the login method. Possible values include `http`, `ontapi`, and `ssh`.
 - `-authentication-method` – The authentication method used for login. Possible values include the following:
 - `domain` – for Active Directory authentication
 - `password` – for password authentication
 - `publickey` – for public-key authentication
 - `-role` – The access-control role name for the login method. At the file system-level, the only role that can be specified is `-role fsxadmin`.
5. The following example demonstrates how to SSH into your file system with your Active Directory credentials if you choose `ssh` for the `-application` type. The `username` is in the format "`domain-name\user-name`", which is the domain name and the username that you provided when creating the account, separated by a backslash and enclosed in quotations.

```
Fsx0123456: :> ssh "CORP\user"@management.fs-abcdef01234567892.fsx.us-east-2.aws.com
```

When prompted to enter a password, use the Active Directory user's password.

Note

The SVM that's used for tunneling needs to have CIFS enabled or be joined to an Active Directory. If you are not enabling CIFS and only joining the tunnel SVM to an Active Directory, ensure that your SVM is joined to your Active Directory. For more information, see [Joining SVMs to a Microsoft Active Directory](#).

Configuring public key authentication

To enable SSH public key authentication, you must first generate an SSH key and associate it with an administrator account by using the `security login publickey create` command. This allows the account to access the SVM. The `security login publickey create` command accepts the following parameters.

Parameter	Description
<code>-vserver</code> (Optional)	The name of the SVM that the account accesses.
<code>-username</code>	The username of the account. The default value, <code>admin</code> , is the default name of the cluster administrator.
<code>-index</code>	The index number of the public key. The default value is 0 if the key is the first key that's created for the account. Otherwise, the default value is one more than the highest existing index number for the account.
<code>-publickey</code>	The OpenSSH public key. Enclose the key in double quotation marks.

Parameter	Description
-role	The access control role that's assigned to the account.
-comment (Optional)	Descriptive text for the public key. Enclose the text in double quotation marks.

The following example associates a public key with the SVM administrator account `svmadmin` for the SVM `svm01`. The public key is assigned index number 5.

```
Fsx0123456::> security login publickey create -vserver svm01 -username svmadmin
-index 5 -publickey "ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEAsPH64CYbUsDQCdW22JnK6J/
vU9upnKzd2zAk9C1f7YaWRUAFNs2Qe5lUmQ3ldi8AD0Vfbr5T6HZPCixNAIzaFciDy7hgnmdj9eNGedGr/
JNrftQbLD1hZybX
+72DpQB0tYWBhe6eDJ1oPLobZBGfMLPXh8VjeU44i7W4+s0hG0E=tsmith@publickey.example.com"
```

Important

You must be an SVM or file system administrator to perform this task.

Migrating to Amazon FSx for NetApp ONTAP

The following sections provide information on how to migrate your existing NetApp ONTAP file systems to Amazon FSx for NetApp ONTAP.

Note

If you plan to use the A11 tiering policy to migrate your data to the capacity pool tier, keep in mind that file metadata is always stored on the SSD tier, and that all new user data is first written to the SSD tier. When data is written to the SSD tier, the background tiering process will begin tiering your data to capacity pool storage, but the tiering process is not immediate and consumes network resources. You need to size your SSD tier to account for file metadata (3-7% of the size of user data), as a buffer for user data before it is tiered to capacity pool storage. We recommend that you do not exceed 80% utilization of your SSD tier.

While migrating data, be sure to monitor your SSD tier using [CloudWatch File system metrics](#) to ensure that it is not filling faster than the tiering process can move data to the capacity pool storage.

Topics

- [Migrating to FSx for ONTAP using NetApp SnapMirror](#)
- [Migrating to FSx for ONTAP using Amazon DataSync](#)

Migrating to FSx for ONTAP using NetApp SnapMirror

You can migrate your NetApp ONTAP file systems to Amazon FSx for NetApp ONTAP using NetApp SnapMirror.

NetApp SnapMirror employs block-level replication between two ONTAP file systems, replicating data from a specified source volume to a destination volume. We recommend using SnapMirror to migrate on-premise NetApp ONTAP file systems to FSx for ONTAP. NetApp SnapMirror's block-level replication is quick and efficient even for file systems with:

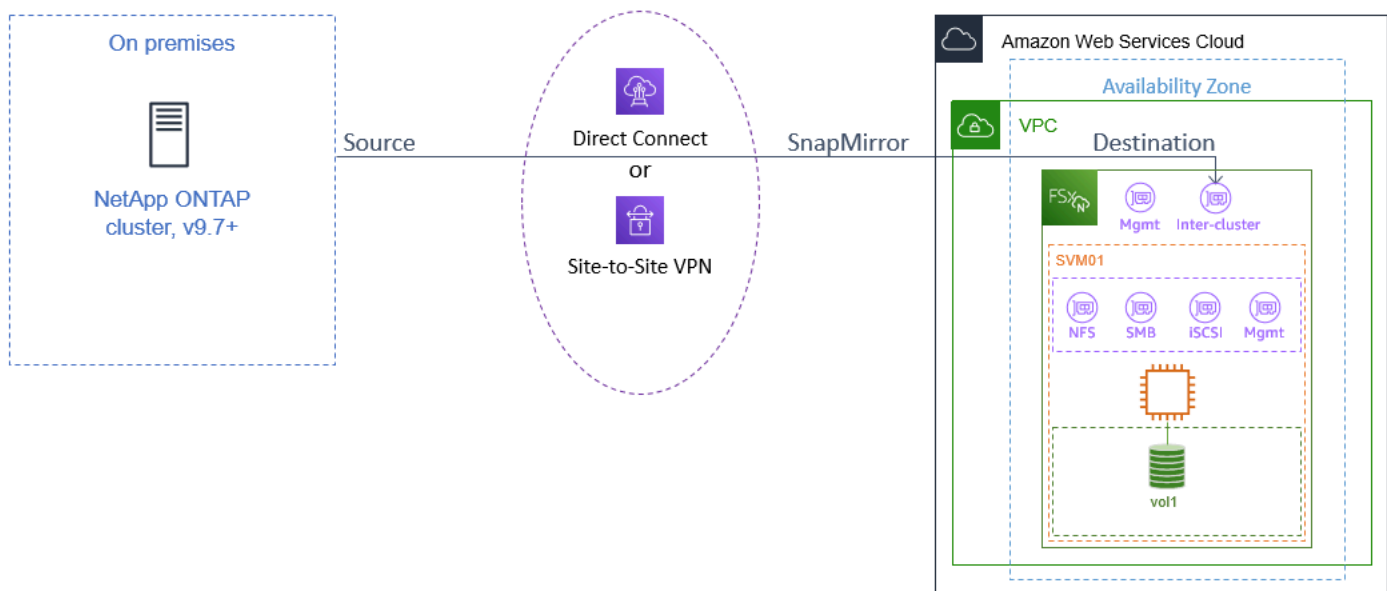
- Complex directory structures
- Over 50 million files

- Very small file sizes (on the order of kilobytes)

When you use SnapMirror to migrate to FSx for ONTAP, deduplicated and compressed data remains in those states, which reduces transfer times and reduces the amount of bandwidth required for migration. Snapshots that exist on the source ONTAP volumes are preserved when migrated to the destination volumes. Migrating your on-premises NetApp ONTAP file systems to FSx for ONTAP involves the following high level tasks:

1. Create the destination volume in Amazon FSx.
2. Gather source and destination logical interfaces (LIFs).
3. Establish cluster peering between the source and destination file systems.
4. Create an SVM peering relationship.
5. Create the SnapMirror relationship.
6. Maintain an updated destination cluster.
7. Cut over to your FSx for ONTAP file system.

The following diagram illustrates the migration scenario described in this section.



Topics

- [Before you begin](#)
- [Create the destination volume](#)
- [Record the source and destination inter-cluster LIFs](#)

- [Establish cluster peering between source and destination](#)
- [Create an SVM peering relationship](#)
- [Create the SnapMirror relationship](#)
- [Transfer data to your FSx for ONTAP file system](#)
- [Cutting over to Amazon FSx](#)


Before you begin

Before you begin using the procedures described in the following sections, be sure that you have met the following prerequisites:

- FSx for ONTAP prioritizes client traffic over background tasks including data tiering, storage efficiency, and backups. When migrating data, and as a general best practice, we recommend that you monitor your SSD tier's capacity to ensure it is not exceeding 80% utilization. You can monitor your SSD tier's utilization using [CloudWatch File system metrics](#). For more information, see [Volume metrics](#).
- If you set the destination volume's data tiering policy to All when migrating your data, all file metadata is stored on the primary SSD storage tier. File metadata is always stored on the SSD-based primary tier, regardless of the volume's data tiering policy. We recommend that you assume a ratio of 1 : 10 for primary tier : capacity pool tier storage capacity.
- The source and destination file systems are connected in the same VPC, or are in networks that are peered using Amazon VPC Peering, Transit Gateway, Amazon Direct Connect or Amazon VPN. For more information, see [Accessing data from within Amazon](#) and [What is VPC peering?](#) in the *Amazon VPC Peering Guide*.
- The VPC security group for the FSx for ONTAP file system has inbound and outbound rules allowing ICMP as well as TCP on ports 443, 10000, 11104, and 11105 for your inter-cluster endpoints (LIFs).
- Verify that the source and destination volumes are running compatible NetApp ONTAP versions before creating a SnapMirror data protection relationship. For more information, see [Compatible ONTAP versions for SnapMirror relationships](#) in NetApp's ONTAP user documentation. The procedures presented here use an on-premise NetApp ONTAP file system for the source.
- Your on-premises (source) NetApp ONTAP file system includes a SnapMirror license.
- You have created a destination FSx for ONTAP file system with an SVM, but you have not created a destination volume. For more information, see [Creating FSx for ONTAP file systems](#).

The commands in these procedures use the following cluster, SVM, and volume aliases:

- *FSx-Dest* – the destination (FSx) cluster's ID (in the format FSxIdabcdef1234567890a).
- *OnPrem-Source* – the source cluster's ID.
- *DestSVM* – the destination SVM name.
- *SourceSVM* – the source SVM name.
- Both the source and destination volume names are vol1.

 **Note**

An FSx for ONTAP file system is referred to as a cluster in all of the ONTAP CLI commands.

The procedures in this section use the following NetApp ONTAP CLI commands.

- [volume create](#) command
- [cluster](#) commands
- [vserver peer](#) commands
- [snapmirror](#) commands

You will use the NetApp ONTAP CLI to create and manage a SnapMirror configuration on your FSx for ONTAP file system. For more information, see [Using the NetApp ONTAP CLI](#).

Create the destination volume

You can create a data protection (DP) destination volume using the Amazon FSx console, the Amazon CLI, and the Amazon FSx API, in addition to the NetApp ONTAP CLI and REST API. For information about creating a destination volume using the Amazon FSx console and Amazon CLI, see [Creating volumes](#).

In the following procedure, you will use the NetApp ONTAP CLI to create a destination volume on your FSx for ONTAP file system. You will need the `fsxadmin` password and the IP address or DNS name of the file system's management port.

1. Establish an SSH session with the destination file system using user `fsxadmin` and the password that you set when you created the file system.

```
ssh fsxadmin@file-system-management-endpoint-ip-address
```

2. Create a volume on the destination cluster that has a storage capacity that is at least equal to the storage capacity of the source volume. Use `-type DP` to designate it as a destination for a SnapMirror relationship.

If you plan to use data tiering, we recommended that you set `-tiering-policy` to `all`. This ensures that your data is immediately transferred to capacity pool storage and prevents you from running out of capacity on your SSD tier. After migration, you can switch `-tiering-policy` to `auto`.

Note

File metadata is always stored on the SSD-based primary tier, regardless of the volume's data tiering policy.

```
FSx-Dest::> vol create -vserver DestSVM -volume vol1 -aggregate aggr1 -size 1g -  
type DP -tiering-policy all
```

Record the source and destination inter-cluster LIFs

SnapMirror uses inter-cluster logical interfaces (LIFs), each with a unique IP address, to facilitate data transfer between source and destination clusters.

1. For the destination FSx for ONTAP file systems, you can retrieve the **Inter-cluster endpoint - IP addresses** from the Amazon FSx console by navigating to the **Administration** tab on your file system's details page.
2. For the source NetApp ONTAP cluster, retrieve the inter-cluster LIF IP addresses using the ONTAP CLI. Run the following command:

```
OnPrem-Source::> network interface show -role intercluster
```

```
Logical          Network
Vserver         Interface  Status   Address/Mask
-----
FSx-Dest
```

```
inter_1    up/up  10.0.0.36/24
inter_2    up/up  10.0.1.69/24
```

Note

For scale-out file systems, there are two inter-cluster IP addresses for each high-availability (HA) pair. Save these values for later.

Save the `inter_1` and `inter_2` IP addresses. They are referenced in the FSx-Dest as `dest_inter_1` and `dest_inter_2` and for OnPrem-Source as `source_inter_1` and `source_inter_2`.

Establish cluster peering between source and destination

Establish a cluster peer relationship on the destination cluster by providing the inter-cluster IP addresses. You will also need to create a passphrase which you will need to enter in when you establish cluster peering on the source cluster.

1. Set up peering on the destination cluster using the following command. For scale-out file systems, you'll need to provide each inter-cluster IP address.

```
FSx-Dest::> cluster peer create -address-family ipv4 -peer-  
addr source_inter_1,source_inter_2
```

Enter the passphrase:

Confirm the passphrase:

Notice: Now use the same passphrase in the "cluster peer create" command in the other cluster.

2. Next, establish the cluster peer relationship on the source cluster. You'll need to enter the passphrase you created above to authenticate. For scale-out file systems, you'll need to provide each inter-cluster IP address.

```
OnPrem-Source::> cluster peer create -address-family ipv4 -peer-  
addr dest_inter_1,dest_inter_2
```

Enter the passphrase:

Confirm the passphrase:

3. Verify the peering was successful using the following command on the source cluster. In the output, Availability should be set to Available.

```
OnPrem-Source::> cluster peer show

Peer Cluster Name  Availability  Authentication
-----
FSx-Dest           Available    ok
```

Create an SVM peering relationship

With cluster peering established, the next step is peering the SVMs. Create an SVM peering relationship on the destination cluster (FSx-Dest) using the `vserver peer` command. Additional aliases used in the following commands are as follows:

- `DestLocalName` – this is name used to identify the destination SVM when configuring SVM peering on the source SVM.
- `SourceLocalName` – this is the name used to identify the source SVM when configuring SVM peering on the destination SVM.

1. Use the following command to create an SVM peering relationship between the source and destination SVMs.

```
FSx-Dest::> vserver peer create -vserver DestSVM -peer-vserver SourceSVM -peer-  
cluster OnPrem-Source -applications snapmirror -local-name SourceLocalName
```

```
Info: [Job 207] 'vserver peer create' job queued
```

2. Accept the peering relationship on the source cluster:

```
OnPrem-Source::> vserver peer accept -vserver SourceSVM -peer-vserver DestSVM -  
local-name DestLocalName
```

```
Info: [Job 211] 'vserver peer accept' job queued
```

3. Verify the SVM peering status using the following command; Peer State should be set to peered in the response.


```
OnPrem-Source::> vserver peer show
```

Peer	Peer	Peer	Peering	Remote	
vserver	Vserver	State	Cluster	Applications	Vserver
svm01	destsvm1	peered	FSx-Dest	snapmirror	svm01

Create the SnapMirror relationship

Now that you have peered the source and destination SVMs, the next steps are to create and initialize the SnapMirror relationship on the destination cluster.

Note

Once you create and initialize a SnapMirror relationship, the destination volumes are read-only until the relationship is broken.

- Use the [snapmirror create](#) command to create the SnapMirror relationship on the destination cluster. The `snapmirror create` command must be used from the destination SVM.

You can optionally use `-throttle` to set the maximum bandwidth (in kB/sec) for the SnapMirror relationship.


```
FSx-Dest::> snapmirror create -source-path SourceLocalName:vol1 -destination-path DestSVM:vol1 -vserver DestSVM -throttle unlimited
```

```
Operation succeeded: snapmirror create for the relationship with destination "DestSVM:vol1".
```

Transfer data to your FSx for ONTAP file system

Now that you've created the SnapMirror relationship, you can transfer data to the destination file system.

1. You can transfer data to the destination file system by running the following command on the destination file system.

 **Note**

Once you run this command, SnapMirror begins transferring snapshots of data from the source volume to the destination volume.

```
FSx-Dest::> snapmirror initialize -destination-path DestSVM:vol1 -source-path SourceLocalName:vol1
```

2. If you are migrating data that is being actively used, you'll need to update your destination cluster so that it remains synced with your source cluster. To perform a one-time update to the destination cluster, run the following command.

```
FSx-Dest::> snapmirror update -destination-path DestSVM:vol1
```

3. You can also schedule hourly or daily updates prior to completing the migration and moving your clients to FSx for ONTAP. You can establish a SnapMirror update schedule using the [snapmirror modify](#) command.

```
FSx-Dest::> snapmirror modify -destination-path DestSVM:vol1 -schedule hourly
```

Cutting over to Amazon FSx

To prepare for the cut over to your FSx for ONTAP file system, do the following:

- Disconnect all clients that write to the source cluster.
- Perform a final SnapMirror transfer to ensure there is no data loss when cutting over.
- Break the SnapMirror relationship.
- Connect all clients to your FSx for ONTAP file system.

1. To ensure that all data from the source cluster is transferred to FSx for ONTAP file system, perform a final Snapmirror transfer.

```
FSx-Dest::> snapmirror update -destination-path DestSVM:vol1
```

2. Ensure that the data migration is complete by verifying that `Mirror State` is set to `Snapmirrored`, and `Relationship Status` is set to `Idle`. You also should ensure that the `Last Transfer End Timestamp` date is as expected, as it shows when the last transfer to the destination volume occurred.
3. Run the following command to show the SnapMirror status.

```
FSx-Dest::> snapmirror show -fields state,status,last-transfer-end-timestamp
```

Source Path	Destination Path	Mirror State	Relationship Status	Last Transfer End Timestamp
Svm01:vol1	svm02:DestVol	Snapmirrored	Idle	09/02 09:02:21

4. Disable any future SnapMirror transfers by using the `snapmirror quiesce` command.

```
FSx-Dest::> snapmirror quiesce -destination-path DestSVM:vol1
```

5. Verify that the `Relationship Status` has changed to `Quiesced` using `snapmirror show`.

```
FSx-Dest::> snapmirror show
```

Source Path	Destination Path	Mirror State	Relationship Status
sourcesvm1:vol1	svm01:DestVol	Snapmirrored	Quiesced

6. During migration, the destination volume is read-only. To enable read/write, you need to break the SnapMirror relationship and cut over to your FSx for ONTAP file system. Break the SnapMirror relationship using the following command.

```
FSx-Dest::> snapmirror break -destination-path DestSVM:vol1
```

```
Operation succeeded: snapmirror break for destination "DestSVM:vol1".
```

7. Once the SnapMirror replication has completed and you have broken the SnapMirror relationship, you can mount the volume to make the data available.

```
FSx-Dest::> vol mount -vserver fsx -volume vol1 -junction-path /vol1
```

The volume is now available with the data from the source volume fully migrated to the destination volume. The volume is also available for clients to read and write to it. If you previously set the `tiering-policy` of this volume to `all`, you can change it to `auto` or `snapshot-only` and your data will automatically transition between storage tiers according to access patterns. To make this data accessible to clients and applications, see [Accessing data](#).

Migrating to FSx for ONTAP using Amazon DataSync

We recommend using Amazon DataSync to transfer data between FSx for ONTAP file systems and non-ONTAP file systems, including FSx for Lustre, FSx for OpenZFS, FSx for Windows File Server, Amazon EFS, Amazon S3, and on-premises filers. If you're transferring files between FSx for ONTAP and NetApp ONTAP, we recommend using [NetApp SnapMirror](#). Amazon DataSync is a data transfer service that simplifies, automates, and accelerates moving and replicating data between self-managed storage systems and Amazon storage services over the internet or Amazon Direct Connect. DataSync can transfer your file system data and metadata, such as ownership, timestamps, and access permissions.

You can use DataSync to transfer files between two FSx for ONTAP file systems, and also move data to a file system in a different Amazon Web Services Region or Amazon account. You can also use DataSync with FSx for ONTAP file systems for other tasks. For example, you can perform one-time data migrations, periodically ingest data for distributed workloads, and schedule replication for data protection and recovery.

In DataSync, a *location* is an endpoint for an FSx for ONTAP file system. For information about specific transfer scenarios, see [Working with locations](#) in the *Amazon DataSync User Guide*.

Note

If you plan to use the `All` tiering policy to migrate your data to the capacity pool tier, keep in mind that file metadata is always stored on the SSD tier, and that all new user data is first written to the SSD tier. When data is written to the SSD tier, the background tiering process will begin tiering your data to capacity pool storage, but the tiering process is not immediate and consumes network resources. You need to size your SSD tier to account for

file metadata (3-7% of the size of user data), as a buffer for user data before it is tiered to capacity pool storage. We recommend that you do not exceed 80% SSD utilization. While migrating data, be sure to monitor your SSD tier using [CloudWatch File system metrics](#) to ensure that it is not filling faster than the tiering process can move data to the capacity pool storage. You can also throttle DataSync transfers to a rate that is lower than the rate that tiering is occurring to ensure that your SSD tier does not exceed 80% utilization. For example, for file systems with a throughput capacity of at least 512 MBps, a 200 MBps throttle will typically balance out the data transfer and data tiering rates.

Prerequisites

To migrate data into your FSx for ONTAP setup, you need a server and network that meet the DataSync requirements. To learn more, see [Requirements for DataSync](#) in the *Amazon DataSync User Guide*.

Basic steps for migrating files using DataSync

Transferring files from a source to a destination using DataSync involves the following basic steps:

- Download and deploy an agent in your environment and activate it (not required if transferring between Amazon Web Services).
- Create a source and destination location.
- Create a task.
- Run the task to transfer files from the source to the destination.

For more information, see the following topics in the *Amazon DataSync User Guide*:

- [Data transfer between self-managed storage and Amazon](#)
- [Creating a location for Amazon FSx for NetApp ONTAP](#)

Monitoring Amazon FSx for NetApp ONTAP

You can use the following services and tools to monitor Amazon FSx for NetApp ONTAP usage and activity:

- **Amazon CloudWatch** – You can monitor file systems using Amazon CloudWatch, which automatically collects and processes raw data from FSx for ONTAP into readable metrics. These statistics are retained for a period of 15 months so that you can access historical information and see how your file system is performing. You can also set alarms based on your metrics over a specified time period and perform one or more actions based on the value of the metrics relative to thresholds that you specify.
- **ONTAP EMS events** – You can monitor your FSx for ONTAP file system by using events generated by ONTAP's Events Management System (EMS). EMS events are notifications of occurrences in your file system, such as iSCSI LUN creation or automatic sizing of volumes.
- **NetApp Cloud Insights** – You can monitor configuration, capacity, and performance metrics for your FSx for ONTAP file systems using the NetApp Cloud Insights service. You can also create alerts based on metric conditions.
- **NetApp Harvest and NetApp Grafana** – You can monitor your FSx for ONTAP file system by using NetApp Harvest and NetApp Grafana. NetApp Harvest monitors ONTAP file systems by collecting performance, capacity, and hardware metrics from FSx for ONTAP file systems. Grafana provides a dashboard where the collected Harvest metrics can be displayed.
- **Amazon CloudTrail** – You can use Amazon CloudTrail to capture all API calls for Amazon FSx as events. These events provide a record of actions taken by a user, role, or Amazon service in Amazon FSx.

Topics

- [Monitoring with Amazon CloudWatch](#)
- [Monitoring FSx for ONTAP workload balance](#)
- [Monitoring FSx for ONTAP EMS events](#)
- [Monitoring with Cloud Insights](#)
- [Monitoring FSx for ONTAP file systems using Harvest and Grafana](#)
- [Logging FSx for ONTAP API Calls with Amazon CloudTrail](#)

Monitoring with Amazon CloudWatch

You can monitor file systems using Amazon CloudWatch, which collects and processes raw data from Amazon FSx for NetApp ONTAP into readable, near real-time metrics. These statistics are retained for a period of 15 months, so that you can access historical information to determine how your file system is performing. FSx for ONTAP metric data is automatically sent to CloudWatch at 1-minute periods by default. For more information about CloudWatch, see [What is Amazon CloudWatch?](#) in the *Amazon CloudWatch User Guide*.

Note

By default, FSx for ONTAP sends metric data to CloudWatch at 1-minute periods except for the following metrics that are sent in 5-minute intervals:

- `FileServerDiskThroughputBalance`
- `FileServerDiskIopsBalance`

CloudWatch metrics for FSx for ONTAP are organized into four categories, which are defined by the dimensions that are used to query each metric. For more information about dimensions, see [Dimensions](#) in the *Amazon CloudWatch User Guide*.

- **File system metrics:** File-system-level performance and storage capacity metrics.
- **Detailed file system metrics:** File-system-level storage metrics per storage tier (SSD and capacity pool).
- **Volume metrics:** Per-volume performance and storage capacity metrics.
- **Detailed volume metrics:** Per-volume storage capacity metrics by storage tier or by the type of data (user, snapshot, or other).

All CloudWatch metrics for FSx for ONTAP are published to the AWS/FSx namespace in CloudWatch.

Topics

- [How to use FSx for ONTAP CloudWatch metrics](#)
- [Accessing CloudWatch metrics](#)
- [File system metrics](#)

- [Scale-out file system metrics](#)
- [Volume metrics](#)
- [Performance warnings and recommendations](#)
- [Creating Amazon CloudWatch alarms to monitor Amazon FSx](#)

How to use FSx for ONTAP CloudWatch metrics

The CloudWatch metrics reported by Amazon FSx provide valuable information about your FSx for ONTAP file systems and volumes.

Topics

- [Monitoring file system metrics in the Amazon FSx console](#)
- [Monitoring volume metrics in the Amazon FSx console](#)

Monitoring file system metrics in the Amazon FSx console

You can use the **Monitoring & performance** panel on your file system's dashboard in the Amazon FSx console to view the metrics that are described in the following table. For more information, see [Accessing CloudWatch metrics](#).

Monitoring & performance	How do I...	Chart	Relevant metrics
Summary	...determine the amount of available storage capacity on my file system?	Available primary storage capacity (bytes)	StorageCapacity {SSD} - StorageUsed {SSD}
	...determine my file system's total client throughput?	Total client throughput (bytes/sec)	$\text{SUM}(\text{DataReadBytes} + \text{DataWriteBytes}) / \text{PERIOD}$ (in seconds)

Monitoring & performance	How do I...	Chart	Relevant metrics
	...determine my file system's total client IOPS?	Total client IOPS (operations/sec)	$\text{SUM}(\text{DataReadOperations} + \text{DataWriteOperations} + \text{MetadataOperations}) / \text{PERIOD (in seconds)}$
	...determine the average latency for the read, write, and metadata operations of my file system?	Average latency (ms/operation)	<p>Average read latency: $\text{DataReadOperationTime} * 1000 / \text{DataReadOperations}$</p> <p>Average write latency: $\text{DataWriteOperationTime} * 1000 / \text{DataWriteOperations}$</p> <p>Average metadata latency: $\text{MetadataOperationTime} * 1000 / \text{MetadataOperations}$</p>
	...determine the distribution of used and free storage capacity on my file system?	Storage distribution	<p>Primary tier available: $\text{StorageCapacity} \{ \text{SSD} \} - \text{StorageUsed} \{ \text{SSD} \}$</p> <p>Primary tier used: $\text{StorageUsed} \{ \text{SSD} \}$</p> <p>Capacity pool used: $\text{StorageUsed} \{ \text{StandardCapacityPool} \}$</p>

Monitoring & performance	How do I...	Chart	Relevant metrics
	...determine the savings from storage efficiencies (compression, deduplication, and compaction)?	Storage efficiency savings	StorageEfficiencySavings
Storage	...determine how much primary storage is available?	Available primary storage capacity (bytes)	StorageCapacity {SSD} - StorageUsed {SSD}
	...determine the percent of used primary storage for my file system?	Primary storage capacity utilization (percent)	StorageCapacity {SSD} * 100/StorageUsed {SSD}
File server performance	...determine if my file system is approaching its network throughput limit?	Network throughput utilization (percent)	NetworkThroughputUtilization
	...determine if my file system is approaching its disk throughput limit?	Disk throughput utilization (percent)	FileServerDiskThroughputUtilization

Monitoring & performance	How do I...	Chart	Relevant metrics
	...determine if my file system has exhausted its allowed burst credits for disk throughput?	Disk throughput – burst balance (percent)	FileServerDiskThroughputBalance
	...determine if my file system is approaching its file servers' SSD IOPS limit?	Disk IOPS – utilization (percent)	FileServerDiskIopsUtilization
	...determine if my file system has exhausted its file servers' allowed burst credits for disk SSD IOPS?	Disk IOPS – burst balance (percent)	FileServerDiskIopsBalance
	...determine the average utilization of the file system's CPU?	CPU utilization (percent)	CPUUtilization
	...determine if my workload is making efficient use of my file system's RAM and NVMe read caches?	Cache hit ratio (percent)	FileServerCacheHitRatio
Disk performance	...determine if my file system is approaching its currently provisioned SSD IOPS capacity?	Disk IOPS – utilization (SSD) (percent)	DiskIopsUtilization

Note

We recommend that you maintain an average throughput capacity utilization of any performance-related dimensions such as network utilization, CPU utilization, and SSD IOPS utilization to under 50%. This ensures that you have enough spare throughput capacity for unexpected spikes in your workload, as well as for any background storage operations (such as storage synchronization, data tiering, or backups).

Monitoring volume metrics in the Amazon FSx console

You can view the **Monitoring** panel on your volume's dashboard in the Amazon FSx console to see additional performance metrics. For more information, see [Accessing CloudWatch metrics](#).

Monitor	How do I...	Chart	Relevant metrics
g	...determine my volume's available storage capacity?	Available storage capacity	StorageCapacity
	...determine my volume's total client throughput?	Total client throughput (bytes/sec)	$\text{SUM}(\text{DataReadBytes} + \text{DataWriteBytes}) / \text{PERIOD}$ (in seconds)
	...determine my volume's total client IOPS?	Total client IOPS (operations/sec)	$\text{SUM}(\text{DataReadOperations} + \text{DataWriteOperations} + \text{MetadataOperations}) / \text{PERIOD}$ (in seconds)
	...determine how many read and write operations are coming from or going to the capacity pool tier?	Capacity Pool IOPS (operations/sec)	Read operations: CapacityPoolReadOperations

Monitor	How do I...	Chart	Relevant metrics
			Write operations: CapacityPoolWriteOperations
	...determine the average latency for the read, write, and metadata operations of my volume?	Average latency (ms/operation)	Average read latency: $\text{DataReadOperationTime} * 1000 / \text{DataReadOperations}$ Average write latency: $\text{DataWriteOperationTime} * 1000 / \text{DataWriteOperations}$ Average metadata latency: $\text{MetadataOperationTime} * 1000 / \text{MetadataOperations}$
	...determine the amount of files or inodes that are available on my volume?	Available files (inodes)	FilesCapacity - FilesUsed
	...determine the distribution of used and free storage capacity on my volume?	Storage distribution	StorageCapacity - StorageUsed

Accessing CloudWatch metrics

You can see Amazon CloudWatch metrics for Amazon FSx in the following ways:

- The Amazon FSx console
- The Amazon CloudWatch console
- The Amazon Command Line Interface (Amazon CLI) for CloudWatch

- The CloudWatch API

The following procedure explains how to view your file system's CloudWatch metrics with the Amazon FSx console.

To view CloudWatch metrics for your file system using the Amazon FSx console

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. In the left navigation pane, choose **File systems**, then choose the file system whose metrics you want to view.
3. On the **Summary** page, choose **Monitoring & performance** from the second panel to view graphs for your file system's metrics.

There are four tabs on the **Monitoring & performance** panel.

- Choose **Summary** (the default tab) to display any active warnings, CloudWatch alarms, and graphs for **File system activity**.
- Choose **Storage** to view storage capacity and utilization metrics.
- Choose **Performance** to view file server and storage performance metrics.
- Choose **CloudWatch alarms** to view graphs of any alarms configured for your file system.

The following procedure explains how to view your volume's CloudWatch metrics with the Amazon FSx console

To view CloudWatch metrics for your volume using the Amazon FSx console

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. In the left navigation pane, choose **Volumes**, then choose the volume whose metrics you want to view.
3. On the **Summary** page, choose **Monitoring** (the default tab) from the second panel to view graphs for your volume's metrics.

The following procedure explains how to view your file system's CloudWatch metrics with the Amazon CloudWatch console.

To view metrics using the Amazon CloudWatch console

1. On the **Summary** page of your file system, choose **Monitoring & performance** from the second panel to view graphs for your file system's metrics.
2. Choose **View in metrics** from the actions menu in the upper right of the graph that you want to view in the Amazon CloudWatch console. This opens the **Metrics** page in the Amazon CloudWatch console.

The following procedure explains how to add FSx for ONTAP file system metrics to a dashboard in the Amazon CloudWatch console.

To add metrics to a Amazon CloudWatch console

1. Choose the set of metrics (**Summary**, **Storage**, or **Performance**) in the **Monitoring & performance** panel of the Amazon FSx console.
2. Choose **Add to dashboard** in the upper right hand of the panel. This opens the Amazon CloudWatch console.
3. Select an existing CloudWatch dashboard from the list, or create a new dashboard. For more information, see [Using Amazon CloudWatch dashboards](#) in the *Amazon CloudWatch User Guide*.

The following procedure explains how to access your file system's metrics with the Amazon CLI.

To access metrics from the Amazon CLI

- Use the CloudWatch [list-metrics](#) CLI command with the `--namespace "AWS/FSx"` parameter. For more information, see the [Amazon CLI Command Reference](#).

The following procedure explains how to access your file system's metrics with the CloudWatch API.

To access metrics from the CloudWatch API

- Call the [GetMetricStatistics](#) API operation. For more information, see the [Amazon CloudWatch API Reference](#).

File system metrics

Your Amazon FSx for NetApp ONTAP file system metrics are classified as either **File system metrics** or **Detailed file system metrics**.

- **File system metrics** are aggregate performance and storage metrics for a single file system that take a single dimension, `FileSystemId`. These metrics measure network performance and storage capacity usage for your file system.
- **Detailed file system metrics** measure your file system's storage capacity and used storage in each storage tier (for example, SSD storage and capacity pool storage). Each metric includes a `FileSystemId`, `StorageTier`, and `DataType` dimension.

Note the following about when Amazon FSx publishes data points for these metrics to CloudWatch:

- For the utilization metrics (any metric whose name ends in *Utilization*, such as `NetworkThroughputUtilization`), there is a data point emitted each period for every active file server or aggregate. For example, Amazon FSx emits one minutely metric per active file server for `FileServerDiskIopsUtilization`, and one minutely metric per aggregate for `DiskIopsUtilization`.
- For all other metrics, there is a single data point emitted each period, corresponding to the total value of the metric across all of your active file servers (such as `DataReadBytes` for file server metrics) or all of your aggregates (such as `DiskReadBytes` for storage metrics).

Topics

- [Network I/O metrics](#)
- [File server metrics](#)
- [Disk I/O metrics](#)
- [Storage capacity metrics](#)
- [Detailed file system metrics](#)

Network I/O metrics

All of these metrics take one dimension, `FileSystemId`.

Metric	Description
NetworkThroughputUtilization	<p>The percent utilization of network throughput for the file system.</p> <p>The Average statistic is the average network throughput utilization of the file system over a specified period.</p> <p>The Minimum statistic is the lowest network throughput utilization of the file system over a specified period.</p> <p>The Maximum statistic is the highest network throughput utilization of the file system over a specified period.</p> <p>Units: Percent</p> <p>Valid statistics: Average, Minimum, and Maximum</p>
NetworkSentBytes	<p>The number of bytes (network I/O) sent by the file system.</p> <p>The Sum statistic is the total number of bytes sent by the file system over a specified period.</p> <p>To calculate sent throughput (bytes per second) for any statistic, divide the statistic by the seconds in the specified period.</p> <p>Units: Bytes</p> <p>Valid statistics: Sum</p>
NetworkReceivedBytes	<p>The number of bytes (network I/O) received by the file system.</p>

Metric	Description
	<p>The Sum statistic is the total number of bytes received by the file system over a specified period.</p> <p>To calculate received throughput (bytes per second) for any statistic, divide the statistic by the seconds in the specified period.</p> <p>Units: Bytes</p> <p>Valid statistics: Sum</p>
DataReadBytes	<p>The number of bytes (network I/O) from reads by clients to the file system.</p> <p>The Sum statistic is the total number of bytes associated with read operations during the specified period. To calculate the average throughput (bytes per second) for a period, divide the Sum statistic by the number of seconds in the specified period.</p> <p>Units: Bytes</p> <p>Valid statistics: Sum</p>

Metric	Description
DataWriteBytes	<p>The number of bytes (network I/O) from writes by clients to the file system.</p> <p>The Sum statistic is the total number of bytes associated with write operations during the specified period. To calculate the average throughput (bytes per second) for a period, divide the Sum statistic by the number of seconds in the specified period.</p> <p>Units: Bytes</p> <p>Valid statistics: Sum</p>
DataReadOperations	<p>The count of read operations (network I/O) from reads by clients to the file system.</p> <p>The Sum statistic is the total number of I/O operations that occurred over a specified period. To calculate the average read operations per second for a period, divide the Sum statistic by the number of seconds in the specified period.</p> <p>Units: Count</p> <p>Valid statistics: Sum</p>

Metric	Description
DataWriteOperations	<p>The count of write operations (network I/O) from writes by clients to the file system.</p> <p>The Sum statistic is the total number of I/O operations that occurred over a specified period. To calculate the average write operations per second for a period, divide the Sum statistic by the number of seconds in the specified period.</p> <p>Units: Count</p> <p>Valid statistics: Sum</p>
MetadataOperations	<p>The count of metadata operations (network I/O) by clients to the file system.</p> <p>The Sum statistic is the total number of I/O operations that occurred over a specified period. To calculate the average metadata operations per second for a period, divide the Sum statistic by the number of seconds in the specified period.</p> <p>Units: Count</p> <p>Valid statistics: Sum</p>

Metric	Description
DataReadOperationTime	<p>The sum of total time spent within the file system for read operations (network I/O) from clients accessing data in the file system.</p> <p>The Sum statistic is the total number of seconds spent by read operations during the specified period. To calculate the average read latency for a period, divide the Sum statistic by the Sum of the DataReadOperations metric over the same period.</p> <p>Units: Seconds</p> <p>Valid statistics: Sum</p>
DataWriteOperationTime	<p>The sum of total time spent within the file system for fulfilling write operations (network I/O) from clients accessing data in the file system.</p> <p>The Sum statistic is the total number of seconds spent by write operations during the specified period. To calculate the average write latency for a period, divide the Sum statistic by the Sum of the DataWriteOperations metric over the same period.</p> <p>Units: Seconds</p> <p>Valid statistics: Sum</p>

Metric	Description
CapacityPoolReadBytes	<p>The number of bytes read (network I/O) from the file system's capacity pool tier.</p> <p>To ensure data integrity, ONTAP performs a read operation on the capacity pool immediately after performing a write operation.</p> <p>The Sum statistic is the total number of bytes read from the file system's capacity pool tier over a specified period. To calculate capacity pool bytes per second, divide the Sum statistic by the seconds in a specified period.</p> <p>Units: Bytes</p> <p>Valid statistics: Sum</p>
CapacityPoolReadOperations	<p>The number of read operations (network I/O) from the file system's capacity pool tier. This translates to a capacity pool read request.</p> <p>To ensure data integrity, ONTAP performs a read operation on the capacity pool immediately after performing a write operation.</p> <p>The Sum statistic is the total number of read operations from the file system's capacity pool tier over a specified period. To calculate capacity pool requests per second, divide the Sum statistic by the seconds in a specified period.</p> <p>Units: Count</p> <p>Valid statistics: Sum</p>

Metric	Description
CapacityPoolWriteBytes	<p>The number of bytes written (network I/O) to the file system's capacity pool tier.</p> <p>To ensure data integrity, ONTAP performs a read operation on the capacity pool immediately after performing a write operation.</p> <p>The Sum statistic is the total number of bytes written to the file system's capacity pool tier over a specified period. To calculate capacity pool bytes per second, divide the Sum statistic by the seconds in a specified period.</p> <p>Units: Bytes</p> <p>Valid statistics: Sum</p>
CapacityPoolWriteOperations	<p>The number of write operations (network I/O) to the file system from the capacity pool tier. This translates to a write request.</p> <p>To ensure data integrity, ONTAP performs a read operation on the capacity pool immediately after performing a write operation.</p> <p>The Sum statistic is the total number of write operations to the file system's capacity pool tier over a specified period. To calculate capacity pool requests per second, divide the Sum statistic by the seconds in a specified period.</p> <p>Units: Count</p> <p>Valid statistics: Sum</p>

File server metrics

All of these metrics take one dimension, `FileSystemId`.

Metric	Description
<p><code>CPUUtilization</code></p>	<p>The percent utilization of the file system's CPU resources.</p> <p>The <code>Average</code> statistic is the average CPU utilization of the file system over a specified period.</p> <p>The <code>Minimum</code> statistic is the lowest CPU utilization of the file system over a specified period.</p> <p>The <code>Maximum</code> statistic is the highest CPU utilization of the file system over a specified period.</p> <p>Units: Percent</p> <p>Valid statistics: <code>Average</code>, <code>Minimum</code>, and <code>Maximum</code></p>
<p><code>FileServerDiskThroughputUtilization</code></p>	<p>The disk throughput between your file server and the primary tier, as a percentage of the provisioned limit determined by throughput capacity.</p> <p>The <code>Average</code> statistic is the average percent utilization of the file servers' disk throughput over a specified period.</p> <p>The <code>Minimum</code> statistic is the lowest percent utilization of the file servers' disk throughput over a specified period.</p>

Metric	Description
	<p>The Maximum statistic is the highest utilization of the file servers' disk throughput over a specified period.</p> <p>Units: Percent</p> <p>Valid statistics: Average, Minimum, and Maximum</p>
FileServerDiskThroughputBalance	<p>The percentage of available burst credits for disk throughput between your file server and the primary tier. This is valid for file systems that are provisioned with a throughput capacity of 512 MBps or less.</p> <p>The Average statistic is the average burst balance available over a specified period.</p> <p>The Minimum statistic is the minimum burst balance available over a specified period.</p> <p>The Maximum statistic is the maximum burst balance available over a specified period.</p> <p>Units: Percent</p> <p>Valid statistics: Average, Minimum, and Maximum</p>

Metric	Description
FileServerDiskIopsBalance	<p>The percentage of available burst credits for disk IOPS between your file server and the primary tier. This is valid for file systems that are provisioned with a throughput capacity of 512 MBps or less.</p> <p>The Average statistic is the average burst balance available over a specified period.</p> <p>The Minimum statistic is the minimum burst balance available over a specified period.</p> <p>The Maximum statistic is the maximum burst balance available over a specified period.</p> <p>Units: Percent</p> <p>Valid statistics: Average, Minimum, and Maximum</p>

Metric	Description
FileServerDiskIopsUtilization	<p>The percentage of IOPS utilization of available disk IOPS capacity for your file server.</p> <p>The Average statistic is the average disk IOPS utilization of the file system over a specified period.</p> <p>The Minimum statistic is the minimum disk IOPS utilization of the file system over a specified period.</p> <p>The Maximum statistic is the maximum disk IOPS utilization of the file system over a specified period.</p> <p>Units: Percent</p> <p>Valid statistics: Average, Minimum, and Maximum</p>

Metric	Description
FileServerCacheHitRatio	<p>The percentage of all read requests that are served by data in the file system's RAM and NVMe caches. A higher percentage means that more reads are served by the file system's read caches.</p> <p>Units: Percent</p> <p>The Average statistic is the average cache hit percent for the file system over a specified period.</p> <p>The Minimum statistic is the lowest cache hit percent for the file system over a specified period.</p> <p>The Maximum statistic is the highest cache hit percent for the file system over a specified period.</p> <p>Valid statistics: Average, Minimum, and Maximum</p>

Disk I/O metrics

All of these metrics take one dimension, `FileSystemId`.

Metric	Description
DiskReadBytes	<p>The number of bytes (disk I/O) from any disk reads to the file system's primary tier.</p> <p>The Sum statistic is the total number of bytes read from the file system over a specified period.</p>

Metric	Description
	<p>To calculate read disk throughput (bytes per second) for any statistic, divide the Sum statistic by the seconds in the specified period.</p> <p>Units: Bytes</p> <p>Valid statistics: Sum</p>
DiskWriteBytes	<p>The number of bytes (disk I/O) from any disk writes to the file system's primary tier.</p> <p>The Sum statistic is the total number of bytes written from the file system over a specified period.</p> <p>To calculate write disk throughput (bytes per second) for any statistic, divide Sum the statistic by the seconds in the specified period.</p> <p>Units: Bytes</p> <p>Valid statistics: Sum</p>

Metric	Description
DiskIopsUtilization	<p>The disk IOPS between your file server and storage volumes, as a percentage of the primary's tiers provisioned disk IOPS limit.</p> <p>The Average statistic is the average disk IOPS utilization of the file system over a specified period.</p> <p>The Minimum statistic is the minimum disk IOPS utilization of the file system over a specified period.</p> <p>The Maximum statistic is the maximum disk IOPS utilization of the file system over a specified period.</p> <p>Units: Percent</p> <p>Valid statistics: Average, Minimum, and Maximum</p>
DiskReadOperations	<p>The number of read operations (disk I/O) from the file system's primary tier.</p> <p>The Sum statistic is the total number of read operations from the primary tier over a specified period.</p> <p>Units: Count</p> <p>Valid statistics: Sum</p>

Metric	Description
DiskWriteOperations	<p>The number of write operations (disk I/O) to the file system's primary tier.</p> <p>The Sum statistic is the total number of write operations to the primary tier over a specified period.</p> <p>Units: Count</p> <p>Valid statistics: Sum</p>

Storage capacity metrics

All of these metrics take one dimension, `FileSystemId`.

Metric	Description
StorageEfficiencySavings	<p>The bytes saved from storage efficiency features (compression, deduplication, and compaction).</p> <p>The Average statistic is the average storage efficiency savings over a specified period. To calculate storage efficiency savings as a percentage of all data stored, over a one minute period, divide <code>StorageEfficiencySavings</code> by the sum of <code>StorageEfficiencySavings</code> and the <code>StorageUsed</code> file system metric, using the Sum statistic for <code>StorageUsed</code>.</p> <p>The Minimum statistic is the minimum storage efficiency savings over a specified period.</p> <p>The Maximum statistic is the maximum storage efficiency savings over a specified period.</p>

Metric	Description
	<p>Units: Bytes</p> <p>Valid statistics: Average, Minimum, and Maximum</p>
StorageUsed	<p>The total amount of physical data stored on the file system, on both the primary (SSD) tier and the capacity pool tier. This metric includes savings from storage-efficiency features, such as data compression and deduplication.</p> <p>Units: Bytes</p> <p>Valid statistics: Average, Minimum, and Maximum</p>

Metric	Description
LogicalDataStored	<p>The total amount of logical data stored on the file system, considering both the SSD tier and the capacity pool tier. This metric includes the total logical size of snapshots and FlexClones, but does not include storage efficiency savings achieved through compression, compaction, and deduplication.</p> <p>To compute storage-efficiency savings in bytes, take the Average of StorageUsed over a given period and subtract it from the Average of LogicalDataStored over the same period.</p> <p>To compute storage-efficiency savings as a percentage of total logical data size, take the Average of StorageUsed over a given period and subtract it from the Average of LogicalDataStored over the same period. Then divide the difference by the Average of LogicalDataStored over the same period.</p> <p>Units: Bytes</p> <p>Valid statistics: Average, Minimum, and Maximum</p>

Detailed file system metrics

Detailed file system metrics are detailed storage-utilization metrics for each of your storage tiers. Detailed file system metrics all have the dimensions `FileSystemId`, `StorageTier`, and `DataType`.

- The `StorageTier` dimension indicates the storage tier that the metric measures, with possible values of `SSD` and `StandardCapacityPool`.

- The `DataType` dimension indicates the type of data that the metric measures, with the possible value `All`.

There is a row for each unique combination of a given metric and dimensional key-value pairs, with a description of what that combination measures.

Metric	Description
StorageCapacityUtilization	<p>The storage capacity utilization for each of your file system's aggregates. There is one metric emitted each minute for each of your file system's aggregates.</p> <p>The <code>Average</code> statistic is the average amount of storage capacity utilization for your file system's performance tier over the specified period.</p> <p>The <code>Minimum</code> statistic is the lowest amount of storage capacity utilization for your file system's performance tier over the specified period.</p> <p>The <code>Maximum</code> statistic is the highest amount of storage capacity utilization for your file system's performance tier over the specified period.</p> <p>Units: Bytes</p> <p>Valid statistics: <code>Average</code>, <code>Minimum</code>, and <code>Maximum</code></p>
StorageCapacity	<p>The total storage capacity of the primary (SSD) tier.</p> <p>Units: Bytes</p>

Metric	Description
	Valid statistics: Maximum
StorageUsed	<p>The used physical storage capacity in bytes, specific to the storage tier. This value includes savings from storage-efficiency features, such as data compression and deduplication. Valid dimension values for <code>StorageTier</code> are <code>SSD</code> and <code>StandardCapacityPool</code> , corresponding to the storage tier that this metric measures. This metric also requires the <code>DataType</code> dimension with the value <code>All</code>.</p> <p>The <code>Average</code>, <code>Minimum</code>, and <code>Maximum</code> statistics are per-tier storage consumption in bytes for the given period.</p> <p>To calculate storage capacity utilization of your primary (SSD) storage tier, divide any of these statistics by the <code>Maximum StorageCapacity</code> over the same period, with the <code>StorageTier</code> dimension equal to <code>SSD</code>.</p> <p>To calculate the free storage capacity of your primary (SSD) storage tier in bytes, subtract any of these statistics from the <code>Maximum StorageCapacity</code> over the same period, with the dimension <code>StorageTier</code> equal to <code>SSD</code>.</p> <p>Units: Bytes</p> <p>Valid statistics: <code>Average</code>, <code>Minimum</code>, and <code>Maximum</code></p>

Scale-out file system metrics

The following metrics are provided for FSx for ONTAP file systems with two or more high-availability (HA) pairs. For the metrics, a datapoint is emitted for each HA pair and for each aggregate (for storage utilization metrics).

Note

If you have a file system with multiple HA pairs, you can also use the [single-HA pair file system metrics](#) and the [volume metrics](#).

Topics

- [Network I/O metrics](#)
- [File server metrics](#)
- [Disk I/O metrics](#)
- [Detailed file system metrics](#)

Network I/O metrics

All of these metrics take two dimensions, `FileSystemId` and `FileServer`.

- `FileSystemId` – Your file system's Amazon resource ID.
- `FileServer` – The name of a file server (or *node*) in ONTAP (for example, `FsxD01234567890abcdef-01`). Odd-numbered file servers are preferred file servers (that is, they service traffic unless the file system has failed over to the secondary file server), while even-numbered file servers are secondary file servers (that is, they serve traffic only when their partner is unavailable). Because of this, secondary file servers typically show less utilization than preferred file servers.

Metric	Description
<code>NetworkThroughputUtilization</code>	Network throughput utilization as a percentage of available network throughput for your file system. This metric is equivalent to the maximum of <code>NetworkSentBytes</code> and

Metric	Description
	<p>NetworkReceivedBytes as a percentage of the network throughput capacity of one HA pair for your file system. All traffic is considered in this metric, including background tasks (such as SnapMirror, tiering, and backups). There is one metric emitted each minute for each of your file system's file servers.</p> <p>The Average statistic is the average network throughput utilization for the given file server over the specified period.</p> <p>The Minimum statistic is the lowest network throughput utilization for the given file server over one minute, for the specified period.</p> <p>The Maximum statistic is the highest network throughput utilization for the given file server over one minute, for the specified period.</p> <p>Units: Percent</p> <p>Valid statistics: Average, Minimum, and Maximum</p>

Metric	Description
NetworkSentBytes	<p>The number of bytes (network IO) sent by your file system. All traffic is considered in this metric, including background tasks (such as SnapMirror, tiering, and backups). There is one metric emitted each minute for each of your file system's file servers.</p> <p>The Sum statistic is the total number of bytes sent over the network by the given file server over the specified period.</p> <p>The Average statistic is the average number of bytes sent over the network by the given file server over the specified period.</p> <p>The Minimum statistic is the lowest number of bytes sent over the network by the given file server over the specified period.</p> <p>The Maximum statistic is the highest number of bytes sent over the network by the given file server over the specified period.</p> <p>To calculate sent throughput (bytes per second) for any statistic, divide the statistic by the seconds in the specified period.</p> <p>Units: Bytes</p> <p>Valid statistics: Sum, Average, Minimum, and Maximum</p>

Metric	Description
NetworkReceivedBytes	<p>The number of bytes (network IO) received by your file system. All traffic is considered in this metric, including background tasks (such as SnapMirror, tiering, and backups). There is one metric emitted each minute for each of your file system's file servers.</p> <p>The Sum statistic is the total number of bytes received over the network by the given file server over the specified period.</p> <p>The Average statistic is the average number of bytes received over the network by the given file server each minute over the specified period.</p> <p>The Minimum statistic is the lowest number of bytes received over the network by the given file server each minute over the specified period.</p> <p>The Maximum statistic is the highest number of bytes received over the network by the given file server each minute over the specified period.</p> <p>To calculate received throughput (bytes per second) for any statistic, divide the statistic by the seconds in the period.</p> <p>Units: Bytes</p> <p>Valid statistics: Sum, Average, Minimum, and Maximum</p>

File server metrics

All of these metrics take two dimensions, `FileSystemId` and `FileServer`.

Metric	Description
<p><code>CPUUtilization</code></p>	<p>The percent utilization of the file system's CPU resources. There is one metric emitted each minute for each of your file system's file servers.</p> <p>The <code>Average</code> statistic is the average CPU utilization of the file system over a specified period.</p> <p>The <code>Minimum</code> statistic is the lowest CPU utilization for the given file server over the specified period.</p> <p>The <code>Maximum</code> statistic is the highest CPU utilization for the given file server over the specified period.</p> <p>Units: Percent</p> <p>Valid statistics: <code>Average</code>, <code>Minimum</code>, and <code>Maximum</code></p>
<p><code>FileServerDiskThroughputUtilization</code></p>	<p>The disk throughput between your file server and aggregate, as a percentage of the provisioned limit determined by throughput capacity. All traffic is considered in this metric, including background tasks (such as <code>SnapMirror</code>, tiering, and backups). This metric is equivalent to the sum of <code>DiskReadBytes</code> and <code>DiskWriteBytes</code> as a percentage of the file server's disk throughput capacity of one HA pair for your file system. There is one</p>

Metric	Description
	<p>metric emitted each minute for each of your file system's file servers.</p> <p>The Average statistic is the average file server disk throughput utilization for the given file server over the specified period.</p> <p>The Minimum statistic is the lowest file server disk throughput utilization for the given file server over the specified period.</p> <p>The Maximum statistic is the highest file server disk throughput utilization for the given file server over the specified period.</p> <p>Units: Percent</p> <p>Valid statistics: Average, Minimum, and Maximum</p>

Metric	Description
FileServerDiskIopsUtilization	<p>The IOPS utilization of available disk IOPS capacity for your file server, as a percentage of its disk IOPS limit. This differs from <code>DiskIopsUtilization</code> in that the utilization of disk IOPS out of the maximum that your file server can handle, as opposed to your provisioned disk IOPS. All traffic is considered in this metric, including background tasks (such as SnapMirror, tiering, and backups). There is one metric emitted each minute for each of your file system's file servers.</p> <p>The <code>Average</code> statistic is the average disk IOPS utilization for the given file server over the specified period.</p> <p>The <code>Minimum</code> statistic is the lowest disk IOPS utilization for the given file server over the specified period.</p> <p>The <code>Maximum</code> statistic is the highest disk IOPS utilization for the given file server over the specified period.</p> <p>Units: Percent</p> <p>Valid statistics: <code>Average</code>, <code>Minimum</code>, and <code>Maximum</code></p>

Metric	Description
FileServerCacheHitRatio	<p>The percentage of all read requests which are served by data that resides in your file system's RAM or NVMe caches for each of your HA pairs (for example, the active file server in an HA pair). A higher percentage indicates a higher ratio of cached reads to total reads. All I/O is considered, including background tasks (such as SnapMirror, tiering, and backups). There is one metric emitted each minute for each of your file system's file servers.</p> <p>Units: Percent</p> <p>The Average statistic is the average cache hit ratio for one of your file system's HA pairs over the specified period.</p> <p>The Minimum statistic is the lowest cache hit ratio for one of your file system's HA pairs over the specified period.</p> <p>The Maximum statistic is the highest cache hit ratio for one of your file system's HA pairs over the specified period.</p> <p>Valid statistics: Average, Minimum, and Maximum</p>

Disk I/O metrics

All of these metrics take two dimensions, `FileSystemId` and `Aggregate`.

- `FileSystemId` – Your file system's Amazon resource ID.
- `Aggregate` – Your file system's performance tier consists of multiple storage pools called *aggregates*. There is one aggregate for each HA pair. For example, aggregate `aggr1`

maps to file server FsxId01234567890abcdef-01 (the active file server) and file server FsxId01234567890abcdef-02 (the secondary file server) in an HA pair.

Metric	Description
DiskReadBytes	<p>The number of bytes (disk IO) from any disk reads from this aggregate. All traffic is considered in this metric, including background tasks (such as SnapMirror, tiering, and backups). There is one metric emitted each minute for each of your file system's aggregates.</p> <p>The Sum statistic is the total number of bytes read each minute from the given aggregate over the specified period.</p> <p>The Average statistic is the average number of bytes read each minute from the given aggregate over the specified period.</p> <p>The Minimum statistic is the lowest number of bytes read each minute from the given aggregate over the specified period.</p> <p>The Maximum statistic is the highest number of bytes read each minute from the given aggregate over the specified period.</p> <p>To calculate read disk throughput (bytes per second) for any statistic, divide the statistic by the seconds in the period.</p> <p>Units: Bytes</p> <p>Valid statistics: Sum, Average, Minimum, and Maximum</p>

Metric	Description
DiskWriteBytes	<p>The number of bytes (disk IO) from any disk writes to this aggregate. All traffic is considered in this metric, including background tasks (such as SnapMirror, tiering, and backups). There is one metric emitted each minute for each of your file system's aggregates.</p> <p>The Sum statistic is the total number of bytes written to the given aggregate over the specified period.</p> <p>The Average statistic is the average number of bytes written to the given aggregate each minute over the specified period.</p> <p>The Minimum statistic is the lowest number of bytes written to the given aggregate each minute over the specified period.</p> <p>The Maximum statistic is the highest number of bytes written to the given aggregate each minute over the specified period.</p> <p>To calculate write disk throughput (bytes per second) for any statistic, divide the statistic by the seconds in the specified period.</p> <p>Units: Bytes</p> <p>Valid statistics: Sum, Average, Minimum, and Maximum</p>

Metric	Description
DiskIopsUtilization	<p>The disk IOPS utilization of one aggregate, as a percentage of the aggregate's disk IOPS limit (that is, the file system's total IOPS divided by the number of HA pairs for your file system). This differs from <code>FileServerDiskIopsUtilization</code> in that it is the utilization of provisioned disk IOPS against your provisioned IOPS limit, as opposed to the maximum disk IOPS supported by the file server (that is, dictated by your configured throughput capacity per HA pair). All traffic is considered in this metric, including background tasks (such as <code>SnapMirror</code>, tiering, and backups). There is one metric emitted each minute for each of your file system's aggregates.</p> <p>The <code>Average</code> statistic is the average disk IOPS utilization for the given aggregate over the specified period.</p> <p>The <code>Minimum</code> statistic is the lowest disk IOPS utilization for the given aggregate over the specified period.</p> <p>The <code>Maximum</code> statistic is the highest disk IOPS utilization for the given aggregate over the specified period.</p> <p>Units: Percent</p> <p>Valid statistics: <code>Average</code>, <code>Minimum</code>, and <code>Maximum</code></p>

Metric	Description
DiskReadOperations	<p>The number of read operations (disk IO) to this aggregate. All traffic is considered in this metric, including background tasks (such as SnapMirror, tiering, and backups). There is one metric emitted each minute for each of your file system's aggregates.</p> <p>The Sum statistic is the total number of read operations performed by the given aggregate over the specified period.</p> <p>The Average statistic is the average number of read operations performed each minute by the given aggregate over the specified period.</p> <p>The Minimum statistic is the lowest number of read operations performed each minute by the given aggregate over the specified period.</p> <p>The Maximum statistic is the highest number of read operations performed each minute by the given aggregate over the specified period.</p> <p>To calculate average disk IOPS over the period, use the Average statistic and divide the result by 60 (seconds).</p> <p>Units: Count</p> <p>Valid statistics: Sum, Average, Minimum, and Maximum</p>

Metric	Description
DiskWriteOperations	<p>The number of write operations (disk IO) to this aggregate. All traffic is considered in this metric, including background tasks (such as SnapMirror, tiering, and backups). There is one metric emitted each minute for each of your file system's aggregates.</p> <p>The Sum statistic is the total number of write operations performed by the given aggregate over the specified period.</p> <p>The Average statistic is the average number of write operations performed each minute by the given aggregate over the specified period.</p> <p>To calculate average disk IOPS over the period, use the Average statistic and divide the result by 60 (seconds).</p> <p>Units: Count</p> <p>Valid statistics: Sum and Average</p>

Detailed file system metrics

Detailed file system metrics are detailed storage-utilization metrics for each of your storage tiers. Detailed file system metrics have either the `FileSystemId`, `StorageTier`, and `DataType` dimensions, or the `FileSystemId`, `StorageTier`, `DataType`, and `Aggregate` dimensions.

- When the `Aggregate` dimension is not supplied, the metrics are for your entire file system. The `StorageUsed` and `StorageCapacity` metrics have a single data point each minute corresponding to the file system's total consumed storage (per storage tier) and total storage capacity (for the SSD tier). Meanwhile, the `StorageCapacityUtilization` metric emits one metric each minute for each aggregate.
- When the `Aggregate` dimension is supplied, the metrics are for each aggregate.

The meaning of the dimensions are as follows:

- **FileSystemId** – Your file system's Amazon resource ID.
- **Aggregate** – Your file system's performance tier consists of multiple storage pools called *aggregates*. There is one aggregate for each HA pair. For example, aggregate `aggr1` maps to file server `FsxId01234567890abcdef-01` (the active file server) and file server `FsxId01234567890abcdef-02` (the secondary file server) in an HA pair.
- **StorageTier** – Indicates the storage tier that the metric measures, with possible values of `SSD` and `StandardCapacityPool`.
- **DataType** – Indicates the type of data that the metric measures, with the possible value `All`.

There is a row for each unique combination of a given metric and dimensional key-value pairs, with a description of what that combination measures.

Metric	Description
StorageCapacityUtilization	<p>The storage capacity utilization for a given file system aggregate. There is one metric emitted each minute for each of your file system's aggregates.</p> <p>The Average statistic is the average amount of storage capacity utilization for a given aggregate over the specified period.</p> <p>The Minimum statistic is the minimum amount of storage capacity utilization for a given aggregate over the specified period.</p> <p>The Maximum statistic is the maximum amount of storage capacity utilization for a given aggregate over the specified period.</p> <p>Units: Bytes</p> <p>Valid statistics: Average, Minimum, and Maximum</p>

Metric	Description
StorageCapacity	<p>The storage capacity for a given file system aggregate. There is one metric emitted each minute for each of your file system's aggregates.</p> <p>The Average statistic is the average amount of storage capacity for a given aggregate over the specified period.</p> <p>The Minimum statistic is the minimum amount of storage capacity for a given aggregate over the specified period.</p> <p>The Maximum statistic is the maximum amount of storage capacity for a given aggregate over the specified period.</p> <p>Units: Bytes</p> <p>Valid statistics: Average, Minimum, and Maximum</p>

Metric	Description
StorageUsed	<p>The used physical storage capacity in bytes, specific to the storage tier. This value includes savings from storage-efficiency features, such as data compression and deduplication. Valid dimension values for <code>StorageTier</code> are <code>SSD</code> and <code>StandardCapacityPool</code> , corresponding to the storage tier that this metric measures. There is one metric emitted each minute for each of your file system's aggregates.</p> <p>The <code>Average</code> statistic is the average amount of physical storage capacity consumed on the given storage tier by the given aggregate over the specified period.</p> <p>The <code>Minimum</code> statistic is the minimum amount of physical storage capacity consumed on the given storage tier by the given aggregate over the specified period.</p> <p>The <code>Maximum</code> statistic is the maximum amount of physical storage capacity consumed on the given storage tier by the given aggregate over the specified period.</p> <p>Units: Bytes</p> <p>Valid statistics: <code>Average</code>, <code>Minimum</code>, and <code>Maximum</code></p>

Volume metrics

Your Amazon FSx for NetApp ONTAP file system can have one or more volumes that store your data. Each of these volumes has a set of metrics, classified as either **Volume metrics** or **Detailed volume metrics**.

- **Volume metrics** are per-volume performance and storage metrics that take two dimensions, `FileSystemId` and `VolumeId`. `FileSystemId` maps to the file system that the volume belongs to.
- **Detailed volume metrics** are per-storage-tier metrics that measure storage consumption per tier with the `StorageTier` dimension (with possible values of `SSD` and `StandardCapacityPool`) and per data type with the `DataType` dimension (with possible values of `User`, `Snapshot`, and `Other`). These metrics have the `FileSystemId`, `VolumeId`, `StorageTier`, and `DataType` dimensions.

Topics

- [Network I/O metrics](#)
- [Storage capacity metrics](#)
- [Detailed volume metrics](#)

Network I/O metrics

All of these metrics take two dimensions, `FileSystemId` and `VolumeId`.

Metric	Description
<code>DataReadBytes</code>	<p>The number of bytes (network I/O) read from the volume by clients.</p> <p>The Sum statistic is the total number of bytes associated with read operations during the specified period. To calculate the average throughput (bytes per second) for a period, divide the Sum statistic by the number of seconds in the specified period.</p> <p>Units: Bytes</p>

Metric	Description
	Valid statistics: Sum
DataWriteBytes	<p>The number of bytes (network I/O) written to the volume by clients.</p> <p>The Sum statistic is the total number of bytes associated with write operations during the specified period. To calculate the average throughput (bytes per second) for a period, divide the Sum statistic by the number of seconds in the specified period.</p> <p>Units: Bytes</p> <p>Valid statistics: Sum</p>
DataReadOperations	<p>The number of read operations (network I/O) on the volume by clients.</p> <p>The Sum statistic is the total number of read operations during the specified period. To calculate the average read operations per second for a period, divide the Sum statistic by the number of seconds in the specified period.</p> <p>Units: Count</p> <p>Valid statistics: Sum</p>

Metric	Description
DataWriteOperations	<p>The number of write operations (network I/O) on the volume by clients.</p> <p>The Sum statistic is the total number of write operations during the specified period. To calculate the average write operations per second for a period, divide the Sum statistic by the number of seconds in the specified period.</p> <p>Units: Count</p> <p>Valid statistics: Sum</p>
MetadataOperations	<p>The number of I/O operations (network I/O) from metadata activities by clients to the volume.</p> <p>The Sum statistic is the total number of metadata operations during the specified period. To calculate the average metadata operations per second for a period, divide the Sum statistic by the number of seconds in the specified period.</p> <p>Units: Count</p> <p>Valid statistics: Sum</p>

Metric	Description
DataReadOperationTime	<p>The sum of total time spent within the volume for read operations (network I/O) from clients accessing data in the volume.</p> <p>The Sum statistic is the total number of seconds spent by read operations during the specified period. To calculate the average read latency for a period, divide the Sum statistic by the Sum of the DataReadOperations metric over the same period.</p> <p>Units: Seconds</p> <p>Valid statistics: Sum</p>
DataWriteOperationTime	<p>The sum of total time spent within the volume for fulfilling write operations (network I/O) from clients accessing data in the volume.</p> <p>The Sum statistic is the total number of seconds spent by write operations during the specified period. To calculate the average write latency for a period, divide the Sum statistic by the Sum of the DataWriteOperations metric over the same period.</p> <p>Units: Seconds</p> <p>Valid statistics: Sum</p>

Metric	Description
MetadataOperationTime	<p>The sum of total time spent within the volume for fulfilling metadata operations (network I/O) from clients that are accessing data in the volume.</p> <p>The Sum statistic is the total number of seconds spent by read operations during the specified period. To calculate the average latency for a period, divide the Sum statistic by the Sum of the MetadataOperations over the same period.</p> <p>Units: Seconds</p> <p>Valid statistics: Sum</p>
CapacityPoolReadBytes	<p>The number of bytes read (network I/O) from the volume's capacity pool tier.</p> <p>To ensure data integrity, ONTAP performs a read operation on the capacity pool immediately after performing a write operation.</p> <p>The Sum statistic is the total number of bytes read from the volume's capacity pool tier over a specified period. To calculate capacity pool bytes per second, divide the Sum statistic by the seconds in a specified period.</p> <p>Units: Bytes</p> <p>Valid statistics: Sum</p>

Metric	Description
CapacityPoolReadOperations	<p>The number of read operations (network I/O) from the volume's capacity pool tier. This translates to a capacity pool read request.</p> <p>To ensure data integrity, ONTAP performs a read operation on the capacity pool immediately after performing a write operation.</p> <p>The Sum statistic is the total number of read operations from the volume's capacity pool tier over a specified period. To calculate capacity pool requests per second, divide the Sum statistic by the seconds in a specified period.</p> <p>Units: Count</p> <p>Valid statistics: Sum</p>
CapacityPoolWriteBytes	<p>The number of bytes written (network I/O) to the volume's capacity pool tier.</p> <p>To ensure data integrity, ONTAP performs a read operation on the capacity pool immediately after performing a write operation.</p> <p>The Sum statistic is the total number of bytes written to the volume's capacity pool tier over a specified period. To calculate capacity pool bytes per second, divide the Sum statistic by the seconds in a specified period.</p> <p>Units: Bytes</p> <p>Valid statistics: Sum</p>

Metric	Description
CapacityPoolWriteOperations	<p>The number of write operations (network I/O) to the volume from the capacity pool tier. This translates to a write request.</p> <p>To ensure data integrity, ONTAP performs a read operation on the capacity pool immediately after performing a write operation.</p> <p>The Sum statistic is the total number of write operations to the volume's capacity pool tier over a specified period. To calculate capacity pool requests per second, divide the Sum statistic by the seconds in a specified period.</p> <p>Units: Count</p> <p>Valid statistics: Sum</p>

Storage capacity metrics

All of these metrics take two dimensions, `FileSystemId` and `VolumeId`.

Metric	Description
StorageCapacity	<p>The size of the volume in bytes.</p> <p>Units: Bytes</p> <p>Valid statistics: Maximum</p>
StorageUsed	<p>The used logical storage capacity of the volume.</p> <p>Units: Bytes</p>

Metric	Description
StorageCapacityUtilization	<p>Valid statistics: Average, Minimum, and Maximum</p> <p>The storage capacity utilization of the volume.</p> <p>Units: Percent</p> <p>Valid statistics: Average</p>
FilesUsed	<p>The used files (number of files or inodes) on the volume.</p> <p>Units: Count</p> <p>Valid statistics: Average, Minimum, and Maximum</p>
FilesCapacity	<p>The total number of inodes that can be created on the volume.</p> <p>Units: Count</p> <p>Valid statistics: Maximum</p>

Detailed volume metrics

Detailed volume metrics take more dimensions than volume metrics, enabling more granular measurements of your data. All detailed volume metrics have the dimensions `FileSystemId`, `VolumeId`, `StorageTier`, and `DataType`.

- The `StorageTier` dimension indicates the storage tier that the metric measures, with possible values of `All`, `SSD`, and `StandardCapacityPool`.
- The `DataType` dimension indicates the type of data that the metric measures, with possible values of `All`, `User`, `Snapshot`, and `Other`.

The following table defines what the `StorageUsed` metric measures for the listed dimensions.

Metric	Description
StorageUsed	<p>The amount of logical space used, in bytes. This metric measures different types of space consumption depending on the dimensions used with this metric. When setting <code>StorageTier</code> to <code>SSD</code> or <code>StandardCapacityPool</code>, and setting <code>DataType</code> to <code>All</code>, this metric measures the logical space usage for this volume for your SSD and capacity pool tiers, respectively. When setting the <code>DataType</code> dimension to <code>User</code>, <code>Snapshot</code>, or <code>Other</code>, and setting <code>StorageTier</code> to <code>All</code>, this metric measures the logical space usage for each respective type of data. The Snapshot data consumption includes the snapshot reserve, which is 5% of the volume's size by default.</p> <p>Units: Bytes</p> <p>Valid statistics: Average, Minimum, and Maximum</p>
StorageCapacityUtilization	<p>The percentage of the volume's used physical disk space.</p> <p>Units: Percent</p> <p>Valid statistics: Maximum</p>

Performance warnings and recommendations

FSx for ONTAP displays a warning for CloudWatch metrics whenever one of these metrics has approached or crossed a predetermined threshold for multiple consecutive data points. These warnings provide you with actionable recommendations that you can use to optimize your file system's performance.

Warnings are accessible in several areas of the **Monitoring & performance** dashboard. All active or recent Amazon FSx performance warnings and any CloudWatch alarms configured for the file system that are in an ALARM state appear in the **Monitoring & performance** panel in the **Summary** section. The warning also appears in the section of the dashboard where the metric graph is displayed.

You can create CloudWatch alarms for any of the Amazon FSx metrics. For more information, see [Creating Amazon CloudWatch alarms to monitor Amazon FSx](#).

Use performance warnings to improve file system performance

Amazon FSx provides actionable recommendations that you can use to optimize your file system's performance. These recommendations describe how you can address a potential performance bottle neck. You can take the recommended action if you expect the activity to continue, or if it's causing an impact to your file system's performance. Depending on which metric has triggered a warning, you can resolve it by increasing either the file system's throughput capacity or storage capacity, as described in the following table.

Dashboard section	If there's a warning for this metric	Do this
Storage	Primary storage capacity utilization	<p>Increase your file system's primary storage capacity if your file system is not already at the maximum SSD storage capacity. For more information, see Modifying SSD storage capacity and provisioned IOPS.</p> <p>If your file system has multiple HA pairs and your primary storage capacity utilization is only higher for a subset of your file system's aggregates (the storage pools that make up your primary storage tier), then you can also rebalance your workload so that your primary storage capacity utilization is more evenly spread across your file system. For more information on rebalancing your workloads, see Monitoring FSx for ONTAP workload balance.</p>

Dashboard section	If there's a warning for this metric	Do this
File server performance	Network throughput	<p>Increase your file system's throughput capacity if your file system is not already at the maximum throughput capacity. For more information on updating throughput capacity, see How to modify throughput capacity.</p> <p>If your file system has multiple HA pairs and utilization is high for only a subset of file servers, then you can also rebalance your workload so that your workload is more evenly utilizing the performance capabilities of each of your file system's HA pairs. For more information on rebalancing your workloads, see Monitoring FSx for ONTAP workload balance.</p>
	Disk throughput	
	Disk IOPS	
	CPU utilization	
Disk performance	Disk IOPS	<p>Increase SSD IOPS if your file system is not already at the maximum SSD IOPS for your file system's current throughput capacity. For more information on updating your file system's provisioned IOPS, see Modifying SSD storage capacity and provisioned IOPS.</p> <p>If your file system has multiple HA pairs and your disk IOPS utilization is only higher for a subset of your file system's aggregates (the storage pools which make up your primary storage tier), then you can also rebalance your workload so that your disk IOPS are utilized more evenly across your file system. For more information on rebalancing your workloads, see Monitoring FSx for ONTAP workload balance.</p>

For more information about file system performance, see [Amazon FSx for NetApp ONTAP performance](#).

Creating Amazon CloudWatch alarms to monitor Amazon FSx

You can create a CloudWatch alarm that sends an Amazon Simple Notification Service (Amazon SNS) message when the alarm changes state. An alarm watches a single metric over a time period that you specify. If needed, the alarm then performs one or more actions based on the value of the metric relative to a given threshold over a number of time periods. The action is a notification sent to an Amazon SNS topic or an Auto Scaling policy.

Alarms invoke actions for sustained state changes only. CloudWatch alarms don't invoke actions only because they are in a particular state; the state must have changed and been maintained for a specified number of periods. You can create an alarm from the Amazon FSx console or the Amazon CloudWatch console.

The following procedures describe how to create alarms using the Amazon FSx console, Amazon Command Line Interface (Amazon CLI), and API.

To set alarms using the Amazon FSx console

1. Open the Amazon FSx console at <https://console.amazonaws.cn/fsx/>.
2. In the left navigation pane, choose **File systems**, and then choose the file system that you want to create the alarm for.
3. On the **Summary** page, choose **Monitoring & performance** from the second panel.
4. Choose the **CloudWatch alarms** tab.
5. Choose **Create CloudWatch alarm**. You are redirected to the CloudWatch console.
6. Choose **Select metric**.
7. In the **Metrics** section, choose **FSx**.
8. Choose a metric category:
 - **File System Metrics**
 - **Detailed File System Metrics**
 - **Volume Metrics**
 - **Detailed Volume Metrics**
9. Choose the metric you want to set the alarm for, and then choose **Select metric**.
10. In the **Conditions** section, choose the conditions you want for the alarm, and then choose **Next**.

Note

Metrics might not be published during file system maintenance. To prevent unnecessary and misleading alarm condition changes and to configure your alarms so that they are resilient to missing data points, see [Configuring how CloudWatch alarms treat missing data](#) in the *Amazon CloudWatch User Guide*.

11. If you want CloudWatch to send you an email or Amazon SNS notification when the alarm state initiates the action, choose an alarm state for **Alarm state trigger**.

For **Send a notification to the following SNS topic**, choose an option. If you choose **Create topic**, you can set the name and email addresses for a new email subscription list. This list is saved and appears in the field for future alarms. Choose **Next**.

Note

If you use **Create topic** to create a new Amazon SNS topic, the email addresses must be verified before they receive notifications. Emails are sent only when the alarm enters an alarm state. If this alarm state change happens before the email addresses are verified, they don't receive a notification.

12. Fill in the **Alarm name** and **Alarm description** fields, and then choose **Next**.
13. On the **Preview and create** page, review the alarm that you're about to create, and then choose **Create alarm**.

To set alarms using the CloudWatch console

1. Open the CloudWatch console at <https://console.amazonaws.cn/cloudwatch/>.
2. Choose **Create Alarm** to start the **Create Alarm Wizard**.
3. Follow the procedure in **To set alarms using the Amazon FSx console**, beginning with step 6.

To set an alarm using the Amazon CLI

- Call the [put-metric-alarm](#) CLI command. For more information, see the [Amazon CLI Command Reference](#).

To set an alarm using the CloudWatch API

- Call the [PutMetricAlarm](#) API operation. For more information, see the [Amazon CloudWatch API Reference](#).

Monitoring FSx for ONTAP workload balance

If you have a file system with multiple HA pairs, then its performance and throughput is spread across each of your HA pairs. FSx for ONTAP automatically balances your files as they are written to your file system, but in rare cases it's possible that your workload data or I/O can become imbalanced across HA pairs, which can impact your workload's overall performance. You can monitor your workload to ensure that it remains balanced across each of your file system's HA pairs (and their commensurate file servers and *aggregates*—the storage pools which make up your primary storage tier).

Topics

- [Primary storage utilization balance](#)
- [File server and disk performance utilization imbalance](#)
- [Mapping CloudWatch dimensions to ONTAP CLI and REST API resources](#)
- [Rebalancing high-traffic clients](#)
- [Rebalancing highly-utilized volumes](#)

Primary storage utilization balance

Your file system's primary storage capacity is divided evenly among each of your HA pairs in storage pools called aggregates. Each HA pair has one aggregate. We recommend that you maintain an average utilization no higher than 80% for your primary storage tier on an ongoing basis. For file systems with multiple HA pairs, we recommend that you maintain an average utilization of up to 80% for every aggregate.

Maintaining 80% utilization ensures there is free space for new incoming data, and maintains a healthy overhead for maintenance operations which can temporarily claim free space on your aggregates.

If you notice that your aggregates are imbalanced, you can either increase your file system's primary storage capacity (commensurately increasing the storage capacity of each aggregate), or

you can move your volumes between aggregates using the [volume move](#) command in the ONTAP CLI.

File server and disk performance utilization imbalance

Your file system's total performance capabilities (such as the network throughput, file server to disk throughput and IOPS, and disk IOPS) is divided evenly among your file system's HA pairs. We recommend that you maintain an average utilization below 50% (and a maximum peak utilization below 80%) for all performance limits on an ongoing basis—this goes for both the overall utilization of your file system's file server resources across all HA pairs, as well as on a per-file server basis.

If you notice that your file server performance utilization is imbalanced—and the file servers on which your workload is imbalanced have an ongoing utilization of over 80%—you can use the ONTAP CLI and REST API to further diagnose the cause of performance imbalance and remediate it. Following is a table of possible imbalance indicators and next steps for further diagnosis.

If your file system's...	Then...
File server disk throughput or file server disk IOPS are imbalanced	You may be experiencing I/O hotspotting on a subset of HA pairs (a subset of your volumes containing an outsized amount of data being accessed) which can limit your workload's overall performance because it's bottlenecked against a subset of HA pairs. For each highly-utilized file server, check the most-utilized volumes to see which volumes have the most activity within an aggregate. For more information on this procedure, see Rebalancing highly-utilized volumes .
Network throughput is imbalanced, but your file server disk throughput, file server disk IOPS, or disk IOPS are not imbalanced	Your data is evenly-distributed across HA pairs, but your clients are not. For the file servers which have more network throughput utilization than others, check the top clients for each file server, then rebalance those clients by unmounting any volumes from those clients and remounting them using a different endpoint on a different HA pair. For more information on this procedure, see Rebalancing high-traffic clients .

Mapping CloudWatch dimensions to ONTAP CLI and REST API resources

Your scale-out file system has Amazon CloudWatch metrics with the FileServer or Aggregate dimension. In order to further diagnose cases of imbalance, you need to map these dimension values to specific file servers (or *nodes*) and aggregates in the ONTAP CLI or REST API.

- For file servers, each file server name maps to a file server (or node) name in ONTAP (for example, FsxId01234567890abcdef-01). Odd-numbered file servers are preferred file servers (that is, they service traffic unless the file system has failed over to the secondary file server), while even-numbered file servers are secondary file servers (that is, they serve traffic only when their partner is unavailable). Because of this, secondary file servers will typically show less utilization than preferred file servers.
- For aggregates, each aggregate name maps to an aggregate in ONTAP (for example, aggr1). There is one aggregate for every HA pair, meaning aggregate aggr1 is shared by file servers FsxId01234567890abcdef-01 (the active file server) and FsxId01234567890abcdef-02 (the secondary file server) in an HA pair, aggregate aggr2 is shared by file servers FsxId01234567890abcdef-03 and FsxId01234567890abcdef-04, and so on.

You can view the mappings between all aggregates and file servers using the ONTAP CLI.

1. To SSH into the NetApp ONTAP CLI of your file system, follow the steps documented in the [Using the NetApp ONTAP CLI](#) section of the *Amazon FSx for NetApp ONTAP User Guide*.

```
ssh fsxadmin@file-system-management-endpoint-ip-address
```

2. Use the [storage aggregate show](#) command, specifying the `-fields node` parameter.

```
::> storage aggregate show -fields node
aggregate                node
-----
aggr1                    FsxId01234567890abcdef-01
aggr2                    FsxId01234567890abcdef-03
aggr3                    FsxId01234567890abcdef-05
aggr4                    FsxId01234567890abcdef-07
aggr5                    FsxId01234567890abcdef-09
aggr6                    FsxId01234567890abcdef-11
6 entries were displayed.
```

Rebalancing high-traffic clients

If you're experiencing I/O imbalance across file servers (specifically with Network throughput utilization), high I/O clients may be the cause. To identify high-traffic clients, use the ONTAP CLI.

1. To SSH into the NetApp ONTAP CLI of your file system, follow the steps documented in the [Using the NetApp ONTAP CLI](#) section of the *Amazon FSx for NetApp ONTAP User Guide*.

```
ssh fsxadmin@file-system-management-endpoint-ip-address
```

2. To view the highest-traffic clients, use the [statistics top client show](#) ONTAP CLI command. You can optionally specify the `-node` parameter to only view the top clients for a specific file server. If you are diagnosing imbalance for a specific file server, use the `-node` parameter, replacing `node_name` with the name of the file server (for example, `FsxId01234567890abcdef-01`).

You can optionally add the `-interval` parameter, providing the interval over which to measure (in seconds) before each report is output. Increasing the interval (for example, to the maximum 300 seconds) provides a longer-term sample for the amount of traffic driven to each volume. The default is 5 (seconds).

```
::> statistics top client show -node FsxId01234567890abcdef-01 [-interval [5,300]]
```

In the output, the top clients are shown by their IP address and port.

Client	Vserver	Node	*Total Ops	Total (Bps)
172.17.236.53:938	svm01	FsxId01234567890abcdef-01	2143	140443648
172.17.236.160:898	svm02	FsxId01234567890abcdef-01	812	53215232

3. You can rebalance a subset of the listed high-traffic clients to other file servers. To do so, unmount the volume from the client and remount it using the DNS name for the SVM's NFS/SMB endpoint—this returns a random endpoint corresponding to a random HA pair.

We recommend you re-use the DNS name, but you have the option to explicitly choose which HA pair a given client mounts. To guarantee that you are mounting a client to a different endpoint, you can instead specify a different endpoint IP address than the one that

corresponds to node that is experiencing high traffic. You can do so by running the following command:

```
::> network interface show -vserver svm_name -lif nfs_smb_management* -fields
address,curr-node
vserver  lif                address          curr-node
-----
svm01   nfs_smb_management_1  172.31.15.89   FsxId01234567890abcdef-01
svm01   nfs_smb_management_3  172.31.8.112   FsxId01234567890abcdef-03
2 entries were displayed.
```

According to the example output for the `statistics top client show` command, client `172.17.236.53` is driving high traffic to `FsxId01234567890abcdef-01`. The output of the `network interface show` command indicates this is the address `172.31.15.89`. To mount to a different endpoint, select any other address (in this example, the only other address is `172.31.8.112`, corresponding to `FsxId01234567890abcdef-03`).

Rebalancing highly-utilized volumes

If you're experiencing I/O imbalance across your volumes or aggregates, you can rebalance volumes in order to redistribute your I/O traffic across your volumes.

Note

If you're experiencing storage utilization imbalance across your aggregates, there is generally not any performance impact unless the high utilization is coupled with I/O imbalance. While you can move volumes between aggregates to balance storage utilization, we recommend only moving volumes if you are seeing a performance impact, as moving volumes can have adverse impact on performance if you don't also consider the I/O driven to each volume you're considering moving.

1. To SSH into the NetApp ONTAP CLI of your file system, follow the steps documented in the [Using the NetApp ONTAP CLI](#) section of the *Amazon FSx for NetApp ONTAP User Guide*.

```
ssh fsxadmin@file-system-management-endpoint-ip-address
```

2. Use the [statistics volume show](#) ONTAP CLI command to view the highest-traffic volumes for a given aggregate, with the following changes:
 - Replace *aggregate_name* with the aggregate's name (for example, aggr1).
 - You can optionally add the `-interval` parameter, providing the interval over which to measure (in seconds) before each report is output. Increasing the interval (for example, to the maximum 300 seconds) provides a longer-term sample for the amount of traffic driven to each volume. The default is 5 (seconds).

```

::> statistics volume show -aggregate aggregate_name -sort-key total_ops [-interval
[5,300]]

```

Depending on the interval you chose, it can take up to 5 minutes to show data. The command shows all volumes in the aggregate, along with the amount of traffic being driven to each aggregate.

Volume	Vserver	Aggregate	*Total Ops	Read Ops	Write Ops	Other Ops	Read (Bps)	Write (Bps)	Latency (us)
vol1__0007	svm1	aggr1	4078	4078	0	0	267255808	0	1092
vol1__0005	svm1	aggr1	4078	4078	0	0	267255808	0	1086
vol1__0003	svm1	aggr1	4077	4077	0	0	267223040	0	1086
vol1__0001	svm1	aggr1	4077	4077	0	0	267239424	0	1087
vol1__0008	svm1	aggr2	2314	2314	0	0	151650304	0	1112
vol1__0006	svm1	aggr2	2144	2144	0	0	140509184	0	1104
vol1__0002	svm1	aggr2	2183	2183	0	0	143065088	0	1106
vol1__0004	svm1	aggr2	2183	2183	0	0	143065088	0	1103

The volume statistics are shown on a per-constituent basis (for example, vol1__0015 is the 15th constituent for FlexGroup vol1). You can see from the example output, the constituents for aggr1 are more highly-utilized than the constituents for aggr2. To balance traffic between aggregates, you can move the constituent volumes between aggregates so that traffic is more evenly distributed.

3. To move a volume between aggregates, use the [volume move start](#) ONTAP CLI command, replacing the following values:
 - Replace *svm_name* with the name of the SVM hosting the volume you're moving.

- Replace *volume_name* with the name of the volume constituent (for example, `vol1__0001`).
- Replace *aggregate_name* with the name of the destination aggregate for the volume.

Important

Volume movement consumes network and disk resources for the source and destination file servers. As a result, the performance of your workload can be impacted by any in-progress volume moves. In addition, there is a cut-over phase of the volume movement process that temporarily pauses I/O for any traffic to the volume.

```
::> volume move start -vserver svm_name -volume volume_name -
destination aggregate_name -foreground false
[Job 1] Job is queued: Move "vol1__0001" in Vserver "svm01" to aggregate "aggr1".
Use the "volume move show -vserver svm01 -volume vol1__0001" command to view the
status of this operation.
```

To check the status of the volume move operation, use the `volume move show` ONTAP CLI command.

```
::> volume move show -vserver svm_name -volume volume_name
      Vserver Name: svm01
      Volume Name: vol1__0001
      Actual Completion Time: -
      Bytes Remaining: 1.00TB
      Specified Action For Cutover: retry_on_failure
      Specified Cutover Time Window: 30
      Destination Aggregate: aggr2
      Destination Node: FsxId01234567890abcdef-03
      Detailed Status: Transferring data: 12.23GB sent.
      Percentage Complete: 1%
      Move Phase: replicating
      Prior Issues Encountered: -
      Estimated Remaining Duration: 00:40:25
      Replication Throughput: 434.3MB/s
      Duration of Move: 00:00:27
      Source Aggregate: aggr2
      Source Node: FsxId01234567890abcdef-01
```

```
Move State: healthy
```

This command shows the estimated time to complete the move, as one of the information fields. When the operation finishes, the same command will show that the Move Phase field is completed.

You should ensure that each FlexGroup is evenly distributed across your aggregates, ideally with the recommended 8 constituents per aggregate. If you move one constituent volume to another aggregate for an otherwise balanced FlexGroup, you should in turn move another (less-utilized) constituent volume to the source aggregate to maintain balance.

Monitoring FSx for ONTAP EMS events

You can monitor FSx for ONTAP file system events using NetAPP ONTAP's native Events Management System (EMS). You can view these events using the NetApp ONTAP CLI.

Topics

- [Overview of EMS events](#)
- [Viewing EMS events](#)
- [EMS event forwarding to a Syslog server](#)

Overview of EMS events

EMS events are automatically generated notifications that alert you when a predefined condition occurs in your FSx for ONTAP file system. These notifications keep you informed so that you can prevent or correct issues that can lead to larger problems, such as storage virtual machine (SVM) authentication issues or full volumes.

By default, events are logged in the Event Management System log. Using EMS, you can monitor events such as user password changes, a constituent within a FlexGroup approaching full capacity, a Logical Unit Number (LUN) was manually brought online or offline, or a volume automatically resizing.

For more information about ONTAP EMS events, see [ONTAP EMS Reference](#) in the NetApp ONTAP Documentation Center. To display the event categories, use the document's left navigation pane.

Note

Only some ONTAP EMS messages are available for FSx for ONTAP file systems.

EMS event descriptions contain event names, severity, possible causes, log messages, and corrective actions that can help you decide how to respond. For example, a [wafl.vol.autoSize.fail](#) event occurs when automatic sizing of a volume fails. According to the event description, the corrective action is to increase the maximum size of the volume while setting the autosize.

Viewing EMS events

Use the NetApp ONTAP CLI **event log show** command to display the contents of the events log. This command is available if you have the `fsxadmin` role on your file system. The command syntax is as follows:

```
event log show [event_options]
```

The most recent events are listed first. By default, this command displays EMERGENCY, ALERT, and ERROR severity-level events with the following information:

- **Time** – The time of the event.
- **Node** – The node on which the event occurred.
- **Severity** – The severity level of the event. To display NOTICE, INFORMATIONAL, or DEBUG severity-level events, use the `-severity` option.
- **Event** – The event name and message.

To display detailed information about events, use one or more of the event options listed in the following table.

Event option	Description
<code>-detail</code>	Displays additional event information.

Event option	Description
-detailtime	Displays detailed event information in reverse chronological order.
-instance	Displays detailed information about all fields.
-node <i>nodename</i> local	Displays a list of events for the node that you specify. Use this option with -seqnum to display detailed information.
-seqnum <i>sequence_number</i>	Selects the events that match this number in the sequence. Use with -node to display detailed information.

Event option	Description
<code>-time <i>MM/DD/YYYY HH:MM:SS</i></code>	<p>Selects the events that happened at this specific time. Use the format: MM/DD/YYYY HH:MM:SS [+HH:MM]. You can specify a time range by using the . . operator between two time statements.</p> <pre>event log show - time "04/17/2023 05:55:00".."04/17/ 2023 06:10:00"</pre> <p>Comparative time values are relative to the current time when you run the command. The following example shows how to display only events that occurred within the last minute:</p> <pre>event log show -time >1m</pre> <p>The month and date fields of this option are not zero-padded. These fields can be single digits; for example, 4/1/2023 06:45:00.</p>

Event option	Description
<code>-severity <i>sev_level</i></code>	<p>Selects the events that match the <i>sev_level</i> value, which must be one of the following:</p> <ul style="list-style-type: none">• EMERGENCY – Disruption• ALERT – Single point of failure• ERROR – Degradation• NOTICE – Information• INFORMATIONAL – Information• DEBUG – Debug information <p>To display all events, specify severity as follows:</p> <pre>event log show -severity <=DEBUG</pre>

Event option	Description
<code>-ems-severity</code> <i>ems_sev_level</i>	<p>Selects the events that match the <i>ems_sev_level</i> value, which must be one of the following:</p> <ul style="list-style-type: none">• NODE_FAULT – Data corruption is detected or the node is unable to provide client service.• SVC_FAULT – A temporary loss of service—typically a transient software fault—is detected.• NODE_ERROR – A hardware error that's not immediately fatal is detected.• SVC_ERROR – A software error that's not immediately fatal is detected.• WARNING – A high-priority message that doesn't indicate a fault.• NOTICE – A normal-priority message that doesn't indicate a fault.• INFO – A low-priority message that doesn't indicate a fault.• DEBUG – A debugging message.

Event option	Description
	<ul style="list-style-type: none">• VAR – A message with variable severity, selected at runtime. <p>To display all events, specify severity as follows:</p> <pre>event log show -ems-severity <=DEBUG</pre>
<code>-source <i>text</i></code>	Selects the events that match the <i>text</i> value. The source is typically a software module.
<code>-message-name <i>message_name</i></code>	Selects the events that match the <i>message_name</i> value. Message names are descriptive, so filtering output by message name displays messages of a specific type.
<code>-event <i>text</i></code>	Selects the events that match the <i>text</i> value. The event field contains the full text of the event, including any parameters.
<code>-kernel-generation-num <i>integer</i></code>	Selects the events that match the <i>integer</i> value. Only events that come from the kernel have kernel generation numbers.

Event option	Description
-kernel-sequence-num <i>integer</i>	Selects the events that match the <i>integer</i> value. Only events that come from the kernel have kernel sequence numbers.
-action <i>text</i>	Selects the events that match the <i>text</i> value. The action field describes what corrective action, if any, you must take to remedy the situation.
-description <i>text</i>	Selects the events that match the <i>text</i> value. The description field describes why the event happened and what it means.
-filter-name <i>filter_name</i>	Selects the events that match the <i>filter_name</i> value. Only events that are included by existing filters that match this value display.
-fields <i>fieldname</i> ,...	Indicates that the command output also includes the specified field or fields. You can use <code>-fields ?</code> to choose the fields that you want to specify.

To view EMS events

1. To SSH into the NetApp ONTAP CLI of your file system, follow the steps documented in the [Using the NetApp ONTAP CLI](#) section of the *Amazon FSx for NetApp ONTAP User Guide*.

```
ssh fsxadmin@file-system-management-endpoint-ip-address
```

2. Use the `event log show` command to display the contents of the event log.

```
::> event log show
Time                Node                Severity            Event
-----
6/30/2023 13:54:19 node1                NOTICE            vifmgr.portup: A link up event was
received on node node1, port e0a.
6/30/2023 13:54:19 node1                NOTICE            vifmgr.portup: A link up event was
received on node node1, port e0d.
```

For information about the EMS events returned by the `event log show` command, refer to the [ONTAP EMS Reference](#) in the NetApp ONTAP Documentation Center.

EMS event forwarding to a Syslog server

You can configure EMS events to forward notifications to a Syslog server. EMS event forwarding is used for real-time monitoring of your file system to determine and isolate root causes for a wide range of issues. If your environment doesn't already contain a Syslog server for event notifications, you must first create one. DNS must be configured on the file system to resolve the Syslog server name.

To configure EMS events to forward notifications to a Syslog server

1. To SSH into the NetApp ONTAP CLI of your file system, follow the steps documented in the [Using the NetApp ONTAP CLI](#) section of the *Amazon FSx for NetApp ONTAP User Guide*.

```
ssh fsxadmin@file-system-management-endpoint-ip-address
```

2. Use the [event notification destination create](#) command to create an event notification destination of type `syslog`, specifying the following attributes:
 - ***dest_name*** – The name of the notification destination that is to be created (for example, `syslog-ems`). An event notification destination name must be 2 to 64 characters long. Valid characters are the following ASCII characters: A-Z, a-z, 0-9, "_", and "-". The name must start and end with: A-Z, a-z, or 0-9.
 - ***syslog_name*** – The Syslog server host name or IP address that Syslog messages are sent to.

- *transport_protocol* – The protocol used to send the events:
 - `udp-unencrypted` – User Datagram Protocol with no security. This is the default protocol.
 - `tcp-unencrypted` – Transmission Control Protocol with no security.
 - `tcp-encrypted` – Transmission Control Protocol with Transport Layer Security (TLS). When this option is specified, FSx for ONTAP verifies the identity of the destination host by validating its certificate.
- *port_number* – The Syslog server port that Syslog messages are sent to. The default value `syslog-port` parameter depends on the setting for the `syslog-transport` parameter. If `syslog-transport` is set to `tcp-encrypted`, the `syslog-port` default value is 6514. If `syslog-transport` is set to `tcp-unencrypted`, `syslog-port` has the default value 601. Otherwise, the default port is set to 514.

```
::> event notification destination create -name dest_name -syslog syslog_name -  
syslog-transport transport_protocol -syslog-port port_number
```

3. Use the [event notification create](#) command to create a new notification of a set of events defined by an event filter to the notification destination created in the previous step, specifying the following attributes:

- *node_name* – The name of the event filter. Events that are included in the event filter are forwarded to the destinations specified in the `-destinations` parameter.
- *dest_name* – The name of the existing notification destination that the event notifications are sent to.

```
::> event notification create -filter-name filter_name -destinations dest_name
```

4. Use the `event notification destination check` command to generate a test message and verify your setup works. Specify the following attributes with the command:

- *node_name* – The name of the node (for example, `FsxId07353f551e6b557b4-01`).
- *dest_name* – The name of the existing notification destination that the event notifications are sent to.

```
::> set diag
```

```
::*> event notification destination check -node node_name -destination-  
name dest_name
```

Monitoring with Cloud Insights

NetApp Cloud Insights is a NetApp service that you can use to monitor your Amazon FSx for NetApp ONTAP file systems alongside your other NetApp storage solutions. With Cloud Insights, you can monitor configuration, capacity, and performance metrics over time to understand your workload's trends and plan for future performance and storage capacity needs. You can also create alerts based on metric conditions that can integrate with your existing workflows and productivity tools.

Note

Cloud Insights isn't supported for scale-out file systems.

Cloud Insights provides:

- **A breadth of metrics and logs** – Collect configuration, capacity, and performance metrics. Understand how your workload is trending with predefined dashboards, alerts, and reports.
- **User analytics and ransomware protection** – With Cloud Secure and ONTAP snapshots you can audit, detect, stop, and repair incidents of user error and ransomware.
- **SnapMirror reporting** – Understand your SnapMirror relationships and set alerts on replication issues.
- **Capacity planning** – Understand the resource requirements of on-premises workloads to help you migrate your workload to a more efficient FSx for ONTAP configuration. You can also use these insights to plan for when more performance or capacity will be needed for your FSx for ONTAP deployment.

For more information about Cloud Insights, see [NetApp Cloud Insights](#) on *NetApp Cloud Central*.

Monitoring FSx for ONTAP file systems using Harvest and Grafana

NetApp Harvest is an open source tool for gathering performance and capacity metrics from ONTAP systems, and is compatible with FSx for ONTAP. You can use Harvest with Grafana for an open source monitoring solution.

Getting started with Harvest and Grafana

The following section details how you can set up and configure Harvest and Grafana to measure your FSx for ONTAP file system's performance and storage capacity utilization.

You can monitor your Amazon FSx for NetApp ONTAP file system by using Harvest and Grafana. NetApp Harvest monitors ONTAP data centers by collecting performance, capacity, and hardware metrics from FSx for ONTAP file systems. Grafana provides a dashboard where the collected Harvest metrics can be displayed.

Supported Harvest dashboards

Amazon FSx for NetApp ONTAP exposes a different set of metrics than does on-premises NetApp ONTAP. Therefore, only the following out-of-the-box Harvest dashboards tagged with `fsx` are currently supported for use with FSx for ONTAP. Some of the panels in these dashboards may be missing information that is not supported.

- ONTAP: Compliance
- ONTAP: Data Protection Snapshots
- ONTAP: Security
- ONTAP: SVM
- ONTAP: Volume

Amazon CloudFormation template

To get started, you can deploy an Amazon CloudFormation template that automatically launches an Amazon EC2 instance running Harvest and Grafana. As an input to the Amazon CloudFormation template, you specify the `fsxadmin` user and the Amazon FSx management endpoint for the file system which will be added as part of this deployment. After the deployment is completed, you can log in to the Grafana dashboard to monitor your file system.

This solution uses Amazon CloudFormation to automate the deployment of the Harvest and Grafana solution. The template creates an Amazon EC2 Linux instance and installs Harvest and Grafana software. To use this solution, download the [fsx-ontap-harvest-grafana.template](#) Amazon CloudFormation template.

Note

Implementing this solution incurs billing for the associated Amazon services. For more information, see the pricing details pages for those services.

Amazon EC2 instance types

When configuring the template, you provide the Amazon EC2 instance type. NetApp's recommendation for the instance size depends on how many file systems you monitor and the number of metrics you choose to collect. With the default configuration, for each 10 file systems you monitor, NetApp recommends:

- CPU: 2 cores
- Memory: 1 GB
- Disk: 500 MB (mostly used by log files)

Following are some sample configurations and the t3 instance type you might choose.

File systems	CPU	Disk	Instance type
Under 10	2 cores	500 MB	t3.micro
10–40	4 cores	1000 MB	t3.xlarge
40+	8 cores	2000 MB	t3.2xlarge

For more information on Amazon EC2 instance types, see [General purpose instances](#) in the *Amazon EC2 User Guide for Linux Instances*.

Instance port rules

When you set up your Amazon EC2 instance, make sure that ports 3000 and 9090 are open for inbound traffic for the security group that the Amazon EC2 Harvest and Grafana instance is in.

Deployment procedure

The following procedure configures and deploys the Harvest/Grafana solution. It takes about five minutes to deploy. Before you start, you must have an FSx for ONTAP file system running in an Amazon Virtual Private Cloud (Amazon VPC) in your Amazon account, and the parameter information for the template listed below. For more information on creating a file system, see [Creating FSx for ONTAP file systems](#).

To launch the Harvest/Grafana solution stack

1. Download the [fsx-ontap-harvest-grafana.template](#) Amazon CloudFormation template. For more information on creating an Amazon CloudFormation stack, see [Creating a stack on the Amazon CloudFormation console](#) in the *Amazon CloudFormation User Guide*.

Note

By default, this template launches in the US East (N. Virginia) Amazon Region. You must launch this solution in an Amazon Web Services Region where Amazon FSx is available. For more information, see [Amazon FSx endpoints and quotas](#) in the *Amazon Web Services General Reference*.

2. For **Parameters**, review the parameters for the template and modify them for the needs of your file system. This solution uses the following default values.

Parameter	Default	Description
InstanceType	t3.micro	The Amazon EC2 instance type. Following are the t3 instance types. <ul style="list-style-type: none"> • t3.micro • t3.small • t3.medium

Parameter	Default	Description
		<ul style="list-style-type: none"> t3.large t3.xlarge t3.2xlarge <p>For the complete list of allowed Amazon EC2 instance type values for this parameter, see the fsx-ontap-harvest-grafana.t emplate.</p>
KeyPair	No default value	The key pair that is used to access the Amazon EC2 instance.
SecurityGroup	No default value	The Security group ID for the Harvest/Grafana Instance. Ensure Inbound ports 3000 and 9090 are open from the clients you wish to use to access your Grafana dashboard.
Subnet Type	No default value	Specify the subnet type, either public or private. Use a public subnet for resources that must be connected to the internet, and a private subnet for resources that won't be connected to the internet. For more information, see Subnet types in the <i>Amazon VPC User Guide</i> .

Parameter	Default	Description
Subnet	No default value	Specify the same subnet as your Amazon FSx for NetApp ONTAP file system's preferred subnet. You can find the file system's Preferred subnet ID in the Amazon FSx console, in the Network & security tab of the FSx for ONTAP file system details page
LatestLinuxAmild	<code>/aws/service/ami-amazon-linux-latest/amzn2-ami-hvm-x86_64-gp2</code>	The latest version of the Amazon Linux 2 AMI in a given Amazon Web Services Region.
FSxEndPoint	No default value	The file system's Management endpoint IP address. You can find the file system's management endpoint IP address in the Amazon FSx console, in the Administration tab of the FSx for ONTAP file system details page.
SecretName	No default value	Amazon Secrets Manager secret name containing the password for the file system's fsxadmin user. This is the password you provided when you created the file system.

3. Choose **Next**.
4. For **Options**, choose **Next**.

5. For **Review**, review and confirm the settings. You must select the check box acknowledging that the template create IAM resources.
6. Choose **Create** to deploy the stack.

You can view the status of the stack in the Amazon CloudFormation console in the **Status** column. You should see a status of **CREATE_COMPLETE** in about five minutes.

Logging in to Grafana

After the deployment has finished, use your browser to log in to the Grafana dashboard at the IP and port 3000 of the Amazon EC2 instance:

```
http://EC2_instance_IP:3000
```

When prompted, use the Grafana default user name (admin) and password (pass). We recommend that you change your password as soon as you log in.

For more information, see the [NetApp Harvest](#) page on GitHub.

Troubleshooting Harvest and Grafana

If you are encountering any data missing mentioned in Harvest and Grafana dashboards or are having trouble setting up Harvest and Grafana with FSx for ONTAP, check the following topics for a potential solution.

Topics

- [SVM and volume dashboards are blank](#)
- [CloudFormation stack rolled back after timeout](#)

SVM and volume dashboards are blank

If the Amazon CloudFormation stack deployed successfully and can contact Grafana but the SVM and volume dashboards are blank, use the following procedure to troubleshoot your environment. You will need SSH access to the Amazon EC2 instance that Harvest and Grafana is deployed on.

1. SSH into the Amazon EC2 instance that your Harvest and Grafana clients are running on.

```
[~]$ ssh ec2-user@ec2_ip_address
```


2. Use the following command to open the `harvest.yml` file and:

- Verify that an entry was created for your FSx for ONTAP instance as `Cluster-2`.
- Verify that the entries for username and password match your `fsxadmin` credentials.

```
[ec2-user@ip-ec2_ip_address ~]$ sudo cat /home/ec2-user/harvest_install/harvest/harvest.yml
```

3. If the password field is blank, open the file in an editor and update it with the `fsxadmin` password, as follows:

```
[ec2-user@ip-ec2_ip_address ~]$ sudo vi /home/ec2-user/harvest_install/harvest/harvest.yml
```

4. Ensure the `fsxadmin` user credentials are stored in Secrets Manager in the following format for any future deployments, replacing `fsxadmin_password` with your password.

```
{"username" : "fsxadmin", "password" : "fsxadmin_password"}
```

CloudFormation stack rolled back after timeout

If you are unable to deploy the CloudFormation stack successfully and it is rolling back with errors, use the following procedure to resolve the issue. You will need SSH access to the EC2 instance deployed by the CloudFormation stack.

1. Redeploy the CloudFormation stack, making sure that automatic rollback is disabled.
2. SSH into the Amazon EC2 instance that your Harvest and Grafana clients are running on.

```
[~]$ ssh ec2-user@ec2_ip_address
```

3. Verify that the docker containers were successfully started using the following command.

```
[ec2-user@ip-ec2_ip_address ~]$ sudo docker ps
```

In the response you should see five containers as follows:

CONTAINER ID	IMAGE	COMMAND	CREATED
STATUS		PORTS	NAMES

```

6b9b3f2085ef  rahulguptajss/harvest  "bin/poller --config..."  8 minutes ago
Restarting (1) 20 seconds ago          harvest_cluster-2
3cf3e3623fde  rahulguptajss/harvest  "bin/poller --config..."  8 minutes ago  Up
About a minute                          harvest_cluster-1
708f3b7ef6f8  grafana/grafana        "/run.sh"                    8 minutes ago  Up
8 minutes                                0.0.0.0:3000->3000/tcp  harvest_grafana
0febee61cab7  prom/alertmanager      "/bin/alertmanager -..."  8
minutes ago  Up 8 minutes          0.0.0.0:9093->9093/tcp
harvest_prometheus_alertmanager
1706d8cd5a0c  prom/prometheus        "/bin/prometheus --c..."  8 minutes ago  Up
8 minutes                                0.0.0.0:9090->9090/tcp  harvest_prometheus

```

4. If the docker containers are not running, check for failures in the `/var/log/cloud-init-output.log` file as follows.

```

[ec2-user@ip-ec2_ip_address ~]$ sudo cat /var/log/cloud-init-output.log
PLAY [Manage Harvest]
*****

TASK [Gathering Facts] *****
ok: [localhost]

TASK [Verify images] *****
failed: [localhost] (item=prom/prometheus) => {"ansible_loop_var": "item",
"changed": false, "item": "prom/prometheus",
"msg": "Error connecting: Error while fetching server API version: ('Connection
aborted.', ConnectionResetError(104, 'Co
nnection reset by peer'))"}
failed: [localhost] (item=prom/alertmanager) => {"ansible_loop_var": "item",
"changed": false, "item": "prom/alertmanage
r", "msg": "Error connecting: Error while fetching server API version: ('Connection
aborted.', ConnectionResetError(104,
'Connection reset by peer'))"}
failed: [localhost] (item=rahulguptajss/harvest) => {"ansible_loop_var": "item",
"changed": false, "item": "rahulguptajs
s/harvest", "msg": "Error connecting: Error while fetching server API version:
('Connection aborted.', ConnectionResetEr
ror(104, 'Connection reset by peer'))"}
failed: [localhost] (item=grafana/grafana) => {"ansible_loop_var": "item",
"changed": false, "item": "grafana/grafana",
"msg": "Error connecting: Error while fetching server API version: ('Connection
aborted.', ConnectionResetError(104, 'Co
nnection reset by peer'))"}

```

```
PLAY RECAP *****
localhost           : ok=1    changed=0    unreachable=0    failed=1
  skipped=0    rescued=0    ignored=0
```

5. If there are failures, execute the following commands to deploy the Harvest and Grafana containers.

```
[ec2-user@ip-ec2_ip_address ~]$ sudo su
[ec2-user@ip-ec2_ip_address ~]$ cd /home/ec2-user/harvest_install
[ec2-user@ip-ec2_ip_address ~]$ /usr/local/bin/ansible-playbook
manage_harvest.yml
[ec2-user@ip-ec2_ip_address ~]$ /usr/local/bin/ansible-playbook
manage_harvest.yml --tags api
```

6. Validate the containers started successfully by running **sudo docker ps** and connecting to your Harvest and Grafana URL.

Logging FSx for ONTAP API Calls with Amazon CloudTrail

Amazon FSx is integrated with Amazon CloudTrail, a service that provides a record of actions taken by a user, role, or an Amazon service in Amazon FSx. CloudTrail captures all Amazon FSx API calls for Amazon FSx for NetApp ONTAP as events. Captured calls include calls from the Amazon FSx console and from code calls to Amazon FSx API operations.

If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for Amazon FSx. If you don't configure a trail, you can still view the most recent events in the CloudTrail console in **Event history**. Using the information collected by CloudTrail, you can determine the request that was made to Amazon FSx. You can also determine the IP address from which the request was made, who made the request, when it was made, and additional details.

To learn more about CloudTrail, see the [Amazon CloudTrail User Guide](#).

Amazon FSx Information in CloudTrail

CloudTrail is enabled on your Amazon account when you create the account. When API activity occurs in Amazon FSx, that activity is recorded in a CloudTrail event along with other Amazon service events in **Event history**. You can view, search, and download recent events in your Amazon account. For more information, see [Viewing events with CloudTrail Event history](#).

For an ongoing record of events in your Amazon account, including events for Amazon FSx, create a trail. A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket. By default, when you create a trail in the console, the trail applies to all Amazon Regions. The trail logs events from all Amazon Regions in the Amazon partition and delivers the log files to the Amazon S3 bucket that you specify. Additionally, you can configure other Amazon services to further analyze and act upon the event data collected in CloudTrail logs. For more information, see the following topics in the *Amazon CloudTrail User Guide*:

- [Creating a trail for your Amazon Web Services account](#)
- [Amazon service integrations with CloudTrail Logs](#)
- [Configuring Amazon SNS notifications for CloudTrail](#)
- [Receiving CloudTrail log files from multiple regions](#) and [Receiving CloudTrail log files from multiple accounts](#)

All Amazon FSx [API calls](#) are logged by CloudTrail. For example, calls to the `CreateFileSystem` and `TagResource` operations generate entries in the CloudTrail log files.

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root or Amazon Identity and Access Management (IAM) user credentials.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another Amazon service.

For more information, see the [CloudTrail userIdentity element](#) in the *Amazon CloudTrail User Guide*.

Understanding Amazon FSx Log File Entries

A *trail* is a configuration that enables delivery of events as log files to an Amazon S3 bucket that you specify. CloudTrail log files contain one or more log entries. An *event* represents a single request from any source and includes information about the requested action, the date and time of the action, request parameters, and so on. CloudTrail log files aren't an ordered stack trace of the public API calls, so they don't appear in any specific order.

The following example shows a CloudTrail log entry that demonstrates the `TagResource` operation when a tag for a file system is created from the console.

```
{
  "eventVersion": "1.05",
  "userIdentity": {
    "type": "Root",
    "principalId": "111122223333",
    "arn": "arn:aws:sts::111122223333:root",
    "accountId": "111122223333",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "sessionContext": {
      "attributes": {
        "mfaAuthenticated": "false",
        "creationDate": "2018-11-14T22:36:07Z"
      }
    }
  },
  "eventTime": "2018-11-14T22:36:07Z",
  "eventSource": "fsx.amazonaws.com",
  "eventName": "TagResource",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "192.0.2.0",
  "userAgent": "console.amazonaws.com",
  "requestParameters": {
    "resourceARN": "arn:aws:fsx:us-east-1:111122223333:file-system/fs-ab12cd34ef56gh789"
  },
  "responseElements": null,
  "requestID": "aEXAMPLE-abcd-1234-56ef-b4cEXAMPLE51",
  "eventID": "bEXAMPLE-gl12-3f5h-3sh4-ab6EXAMPLE9p",
  "eventType": "AwsApiCall",
  "apiVersion": "2018-03-01",
  "recipientAccountId": "111122223333"
}
```

The following example shows a CloudTrail log entry that demonstrates the `UntagResource` action when a tag for a file system is deleted from the console.

```
{
  "eventVersion": "1.05",
  "userIdentity": {
    "type": "Root",
    "principalId": "111122223333",
    "arn": "arn:aws:sts::111122223333:root",
    "accountId": "111122223333",
```

```
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "sessionContext": {
      "attributes": {
        "mfaAuthenticated": "false",
        "creationDate": "2018-11-14T23:40:54Z"
      }
    }
  },
  "eventTime": "2018-11-14T23:40:54Z",
  "eventSource": "fsx.amazonaws.com",
  "eventName": "UntagResource",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "192.0.2.0",
  "userAgent": "console.amazonaws.com",
  "requestParameters": {
    "resourceARN": "arn:aws:fsx:us-east-1:111122223333:file-system/fs-
ab12cd34ef56gh789"
  },
  "responseElements": null,
  "requestID": "aEXAMPLE-abcd-1234-56ef-b4cEXAMPLE51",
  "eventID": "bEXAMPLE-gl12-3f5h-3sh4-ab6EXAMPLE9p",
  "eventType": "AwsApiCall",
  "apiVersion": "2018-03-01",
  "recipientAccountId": "111122223333"
}
```

Quotas

Following, you can find out about quotas when working with Amazon FSx for NetApp ONTAP.

Topics

- [Quotas that you can increase](#)
- [Resource quotas for each file system](#)

Quotas that you can increase

Following are the quotas for Amazon FSx for NetApp ONTAP for each Amazon Web Services account, per Amazon Web Services Region, that you can increase.

Resource	Default	Description
ONTAP file systems	100	The maximum number of Amazon FSx for NetApp ONTAP file systems that you can create in this account.
ONTAP SSD storage capacity	524,288	The maximum amount of SSD storage capacity (in GiB) for all Amazon FSx for NetApp ONTAP file systems that you can have in this account.
ONTAP throughput capacity	10,240	The maximum amount of throughput capacity (in MBps) for all Amazon FSx for NetApp ONTAP file systems that you can have in this account.
ONTAP SSD IOPS	1,000,000	The maximum amount of SSD IOPS for all Amazon FSx for NetApp ONTAP file systems

Resource	Default	Description
		that you can have in this account.
ONTAP backups per file system	10,000	The maximum number of user-initiated volume backups for all Amazon FSx for NetApp ONTAP file systems that you can have in this account.

To request a quota increase

1. Open the [Amazon Web Services Support](#) page, sign in if necessary, and then choose **Create case**.
2. For **Create case**, choose **Account and billing support**.
3. In the **Case details** panel make the following entries:
 - For **Type** choose **Account**.
 - For **Category** choose **Other Account Issues**.
 - For **Subject** enter **Amazon FSx for NetApp ONTAP service limit increase request**.
 - Provide a detailed **Description** of your request, including:
 - The FSx quota that you want increased, and the value you want it increased to, if known.
 - The reason why you are seeking the quota increase.
 - The file system ID and region for each file system you are requesting an increase for.
4. Provide your preferred **Contact options** and choose **Submit**.

Resource quotas for each file system

The following table lists the quotas on Amazon FSx for NetApp ONTAP resources for each file system in an Amazon Web Services Region.

Resource	Limit per file system
Minimum SSD storage capacity	1,024 GiB per high-availability (HA) pair
Maximum SSD storage capacity	<ul style="list-style-type: none"> Scale-out: 512 TiB per HA pair, up to 1 PiB Scale-up: 192 TiB
Maximum SSD IOPS	<p>Scale-out:</p> <ul style="list-style-type: none"> 200,000 per HA pair (up to 12 pairs) <p>Scale-up:</p> <ul style="list-style-type: none"> 160,000 in US East (Ohio) Region, US East (N. Virginia) Region, US West (Oregon) Region, and Europe (Ireland) 80,000 in all other Amazon Web Services Regions where FSx for ONTAP is available
Minimum throughput capacity	<ul style="list-style-type: none"> Scale-out: 3,072 MBps per HA pair Scale-up: 128 MBps
Maximum throughput capacity	<p>Scale-out:</p> <ul style="list-style-type: none"> 73,728 MBps¹ <p>Scale-up:</p> <ul style="list-style-type: none"> 4,096 MBps² in US East (Ohio) Region, US East (N. Virginia) Region, US West (Oregon) Region, and Europe (Ireland)

Resource	Limit per file system
	<ul style="list-style-type: none"> 2,048 MBps in all other Amazon Web Services Regions where FSx for ONTAP is available
Maximum number of volumes	<ul style="list-style-type: none"> Scale-out: 1,000 Scale-up: 500
Maximum number of snapshots	1,023 per volume ³
Maximum number of backups	4,091 per volume ⁴
Maximum number of SVMs	<p>Scale-out:</p> <ul style="list-style-type: none"> 5 <p>Scale-up:</p> <ul style="list-style-type: none"> 6 (128 MBps throughput capacity) 6 (256 MBps throughput capacity) 14 (512 MBps throughput capacity) 14 (1,024 MBps throughput capacity) 24 (2,048 MBps throughput capacity) 24 (4,096 MBps throughput capacity)
Maximum number of tags	50
Maximum retention period for automated backups	90 days
Maximum retention period for user-initiated backups	No retention limit

Note

¹ On a scale-out file system with 12 HA pairs (6,144 MBps per HA pair). For more information, see [High-availability \(HA\) pairs](#).

² To provision 4 GBps of throughput capacity, your FSx for ONTAP scale-up file system requires a configuration of the maximum SSD IOPS (160,000) and a minimum of 5,120 GiB of SSD storage capacity in a supported Amazon Web Services Region. For more information about which Amazon Web Services Regions support 4,096 MBps of throughput capacity, see [Impact of throughput capacity on performance](#).

³ You can store up to 1,023 snapshots per volume at any point in time. Once you reach this limit, you must delete an existing snapshot before a new snapshot of your volume can be created.

⁴ You can store up to 4,091 backups per volume at any point in time. Once you reach this limit, you must delete an existing backup before a new backup of your volume can be created.

Troubleshooting Amazon FSx for NetApp ONTAP

Use the following sections to help troubleshoot problems that you have with FSx for ONTAP.

Topics

- [My Multi-AZ file system is in a MISCONFIGURED state](#)
- [You can't access your file system](#)
- [You can't join a storage virtual machine \(SVM\) to Active Directory](#)
- [You can't delete a storage virtual machine or volume](#)
- [Automatic daily backups fail due to insufficient volume capacity](#)
- [You have insufficient volume capacity](#)
- [Troubleshooting network issues](#)

My Multi-AZ file system is in a MISCONFIGURED state

There are a number of potential causes for a file system to be in a MISCONFIGURED state, each with their own resolution, as follows.

Topics

- [The VPC owner account has disabled Multi-AZ VPC sharing](#)
- [You can't create a new SVM on a Multi-AZ file system](#)

The VPC owner account has disabled Multi-AZ VPC sharing

Multi-AZ file systems created by a participant Amazon Web Services account in a shared VPC subnet will go into a MISCONFIGURED state for one of the following reasons:

- The owner account that shared the VPC subnet has disabled Multi-AZ VPC sharing support for FSx for ONTAP file systems.
- The owner account has unshared the VPC subnet.

If the owner account has unshared the VPC subnet, you will see the following message in the console for that file system:

The vpc ID *vpc-012345abcde* does not exist

You need to contact the owner account that shared the VPC subnet with you to resolve the issue. For more information see [Creating FSx for ONTAP file systems in shared subnets](#) for more information.

You can't create a new SVM on a Multi-AZ file system

For Multi-AZ file systems created by a participant Amazon Web Services account in a shared VPC, you will be unable to create a new SVM for one of the following reasons:

- The owner account that shared the VPC subnet has disabled Multi-AZ VPC sharing support for FSx for ONTAP file systems.
- The owner account has unshared the VPC subnet.

You need to contact the owner account that shared the VPC subnet with you to resolve the issue. For more information see [Creating FSx for ONTAP file systems in shared subnets](#) for more information.

You can't access your file system

There are a number of potential causes for being unable to access your file system, each with their own resolution, as follows.

Topics

- [The file system's elastic network interface was modified or deleted](#)
- [The Elastic IP address attached to the file system's elastic network interface was deleted](#)
- [The file system's VPC security group lacks the required inbound rules](#)
- [The compute instance's VPC security group lacks the required outbound rules](#)
- [The compute instance's subnet doesn't use any of the route tables associated with your file system](#)
- [Amazon FSx can't update route table for Multi-AZ file systems created using Amazon CloudFormation](#)
- [Can't access a file system over iSCSI from a client in another VPC](#)
- [The owning account has unshared the VPC subnet](#)

- [Can't access a file system over NFS, SMB, the ONTAP CLI, or the ONTAP REST API from a client in another VPC or on-premises](#)

The file system's elastic network interface was modified or deleted

You must not modify or delete any of the file system's elastic network interfaces. Modifying or deleting a network interface can cause a permanent loss of connection between your virtual private cloud (VPC) and your file system. Create a new file system, and don't modify or delete the Amazon FSx network interface. For more information, see [File System Access Control with Amazon VPC](#).

The Elastic IP address attached to the file system's elastic network interface was deleted

Amazon FSx doesn't support accessing file systems from the public Internet. Amazon FSx automatically detaches any Elastic IP address which is a public IP address reachable from the Internet that gets attached to a file system's elastic network interface. For more information, see [Supported clients](#).

The file system's VPC security group lacks the required inbound rules

Review the inbound rules specified in [Amazon VPC security groups](#), and make sure that the security group associated with your file system has the corresponding inbound rules.

The compute instance's VPC security group lacks the required outbound rules

Review the outbound rules specified in [Amazon VPC security groups](#), and make sure that the security group associated with your compute instance has the corresponding outbound rules.

The compute instance's subnet doesn't use any of the route tables associated with your file system

FSx for ONTAP creates endpoints for accessing your file system in a VPC route table. We recommend that you configure your file system to use all of the VPC route tables that are associated with the subnets in which your clients are located. By default, Amazon FSx uses your VPC's main route table. You can optionally specify one or more route tables for Amazon FSx to use when you create your file system.

If you can ping your file system's Intercluster endpoint but cannot ping your file system's Management endpoint (see [File system resources](#) for more information), your client is likely not in a subnet that's associated with one of your file system's route tables. To access your file system, associate one of your file system's route tables with your client's subnet. For information about updating your file system's Amazon VPC route tables, see [Updating a file system](#).

Amazon FSx can't update route table for Multi-AZ file systems created using Amazon CloudFormation

Amazon FSx manages VPC route tables for Multi-AZ file systems using tag-based authentication. These route tables are tagged with `Key: AmazonFSx; Value: ManagedByAmazonFSx`. When creating or updating FSx for ONTAP Multi-AZ file systems using Amazon CloudFormation we recommend that you add the `Key: AmazonFSx; Value: ManagedByAmazonFSx` tag manually.

If you're unable to reach your Multi-AZ file system, check to see if the VPC route tables associated with the file system are tagged with `Key: AmazonFSx; Value: ManagedByAmazonFSx`. If they are not, then Amazon FSx cannot update those route tables to route the floating IP addresses of the management and data ports to the active file server when a failover event occurs. For information about updating your file system's Amazon VPC route tables, see [Updating a file system](#).

Can't access a file system over iSCSI from a client in another VPC

To access a file system over the Internet Small Computer Systems Interface (iSCSI) protocol from a client in another VPC, you can configure Amazon VPC peering or Amazon Transit Gateway between the VPC associated with your file system and the VPC in which your client resides. For more information, see [Create and accept VPC peering connections](#) in the *Amazon Virtual Private Cloud* guide.

The owning account has unshared the VPC subnet

If you created your file system in a VPC subnet that has been shared with you, the owning account may have unshared the VPC subnet.

If the owner account has unshared the VPC subnet, you will see the following message in the console for that file system:

```
The vpc ID vpc-012345abcde does not exist
```

You will need to contact the owning account so that they can re-share the subnet with you.

Can't access a file system over NFS, SMB, the ONTAP CLI, or the ONTAP REST API from a client in another VPC or on-premises

To access a file system over Network File System (NFS), Server Message Block (SMB), or the NetApp ONTAP CLI and REST API from a client in another VPC or on premises, you must configure routing using Amazon Transit Gateway between the VPC associated with your file system and the network in which your client resides. For more information, see [Accessing data](#).

You can't join a storage virtual machine (SVM) to Active Directory

If you're unable to join an SVM to an Active Directory (AD), first review [Joining SVMs to a Microsoft Active Directory](#). Common problems that prevent an SVM from joining to your Active Directory are listed in the following sections, including the error messages generated for each circumstance.

Topics

- [The SVM NetBIOS name is the same as the NetBIOS name for the home domain.](#)
- [The SVM is already joined to another Active Directory](#)
- [Amazon FSx can't connect to your Active Directory domain controllers because the SVM's NetBIOS name is already in use](#)
- [Amazon FSx can't communicate with your Active Directory domain controllers](#)
- [Amazon FSx can't connect to your Active Directory due to unmet port requirements or service account permissions](#)
- [Amazon FSx can't connect to your Active Directory domain controllers because the service account credentials are not valid](#)
- [Amazon FSx can't connect to your Active Directory domain controllers because of insufficient service account credentials](#)
- [Amazon FSx can't communicate with your Active Directory DNS servers or domain controllers](#)
- [Amazon FSx can't communicate with your Active Directory because of a invalid Active Directory domain name.](#)
- [The service account can't access the administrators group specified in the SVM Active Directory configuration](#)

- [Amazon FSx can't connect to the Active Directory domain controllers because the organizational unit specified doesn't exist or isn't accessible](#)

The SVM NetBIOS name is the same as the NetBIOS name for the home domain.

Joining an SVM to your self-managed Active Directory fails with the following error message:

Amazon FSx is unable to establish a connection with your Active Directory. This is because the server name you specified is the NetBIOS name of the home domain. To fix this problem, choose a NetBIOS name for your SVM that is different from the NetBIOS name of the home domain. Then reattempt to join your SVM to your Active Directory.

To resolve this issue, follow the procedure described in [Joining an SVM to an Active Directory using the Amazon Web Services Management Console, Amazon CLI and API](#) to reattempt joining your SVM to your AD. Ensure that you use a NetBIOS name for your SVM that's different than the NetBIOS name of the Active Directory's home domain.

The SVM is already joined to another Active Directory

Joining an SVM to an Active Directory fails with the following error message:

Amazon FSx is unable to establish a connection to your Active Directory. This is because the SVM is already joined to a domain. To join this SVM to a different domain, you can use the ONTAP CLI or REST API to unjoin this SVM from Active Directory. Then reattempt to join your SVM to a different Active Directory.

To resolve the issue, do the following:

1. Use the NetApp ONTAP CLI to unjoin the SVM from its current Active Directory. For more information, see [Unjoin an Active Directory from your SVM using the NetApp ONTAP CLI](#).
2. Follow the procedure described in [Joining an SVM to an Active Directory using the Amazon Web Services Management Console, Amazon CLI and API](#) to reattempt joining your SVM to the new AD.

Amazon FSx can't connect to your Active Directory domain controllers because the SVM's NetBIOS name is already in use

Creating an SVM joined to your self-managed AD fails with the following error message:

Amazon FSx is unable to establish a connection with your Active Directory. This is because the NetBIOS (computer) name you specified is already in-use in your Active Directory. To fix this problem, pick a NetBIOS name for your SVM that is not in use in your Active Directory., specifying a NetBIOS (computer) Then reattempt to join your SVM to your Active Directory.

To resolve this issue, follow the procedure described in [Joining an SVM to an Active Directory using the Amazon Web Services Management Console, Amazon CLI and API](#) to reattempt joining your SVM to your AD. Ensure that you use a NetBIOS name for your SVM that's unique and not already in use in your Active Directory.

Amazon FSx can't communicate with your Active Directory domain controllers

Joining an SVM to your self-managed AD fails with the following error message:

Amazon FSx is unable to communicate with your Active Directory. To fix this problem, ensure that network traffic is allowed between Amazon FSx and your domain controllers. Then reattempt to join your SVM to your Active Directory.

To resolve this issue, do the following:

1. Review the requirements described in [Network configuration requirements](#), and make changes needed to enable network communications between Amazon FSx and your AD.
2. Once Amazon FSx is able to communicate with your AD, follow the procedure described in [Joining an SVM to an Active Directory using the Amazon Web Services Management Console, Amazon CLI and API](#) and reattempt joining your SVM to your AD.

Amazon FSx can't connect to your Active Directory due to unmet port requirements or service account permissions

Joining an SVM to your self-managed AD fails with the following error message:

Amazon FSx is unable to establish a connection with your Active Directory. This is due to either the port requirements for your Active Directory not being met, or the service account provided not having permissions to join the storage virtual machine to the domain with the specified organization unit. To fix this problem, update your storage virtual machine's Active Directory configuration after resolving any permissions issues with ports and service accounts, as recommended in the Amazon FSx user guide.

To resolve this issue, do the following:

1. Review the requirements described in [Network configuration requirements](#), and make changes needed to meet the networking requirements and make sure communications are enabled on the required ports
2. Review the service account requirements described in [Active Directory service account requirements](#). Ensure that service account has the delegated permissions necessary to join your SVM to the AD domain using the specified organizational unit.
3. Once you have made changes to the port permissions or the service account, follow the procedure described in [Joining an SVM to an Active Directory using the Amazon Web Services Management Console, Amazon CLI and API](#) and reattempt joining your SVM to your AD.

Amazon FSx can't connect to your Active Directory domain controllers because the service account credentials are not valid

Joining an SVM to your self-managed Active Directory fails with the following error message:

Amazon FSx is unable to establish a connection with your Active Directory domain controller(s) because the service account credentials provided are invalid. To fix this problem, update your storage virtual machine's Active Directory configuration with a valid service account.

To resolve this issue, use the procedure described in [Updating an existing SVM Active Directory configuration using the Amazon Web Services Management Console, Amazon CLI, and API](#) to update the SVM's service account credentials. When entering the service account user name, be sure to include only the user name (for example, `ServiceAcct`), and don't include any domain prefix (for example, `corp.com\ServiceAcct`) or domain suffix (for example, `ServiceAcct@corp.com`). Don't use the distinguished name (DN) when entering the service account user name (for example, `CN=ServiceAcct,OU=example,DC=corp,DC=com`).

Amazon FSx can't connect to your Active Directory domain controllers because of insufficient service account credentials

Joining an SVM to your self-managed Active Directory fails with the following error message:

Amazon FSx is unable to establish a connection with your Active Directory domain controller(s). This is due to either the port requirements for the Active Directory have not been met, or the service account provided does not have permission to join the storage virtual machine to the domain with the specified organizational unit.

To resolve this issue, make sure that you have delegated the required permissions to the service account that you provided. The service account must be able to create and delete computer objects in the OU in the domain to which you're joining the file system. The service account also needs, at a minimum, to have permissions to do the following:

- Reset passwords
- Restrict accounts from reading and writing data
- Validated ability to write to the DNS hostname
- Validated ability to write to the service principal name
- Ability to create and delete computer objects
- Validated ability to read and write Account Restrictions

For more information about creating a service account with correct permissions, see [Active Directory service account requirements](#) and [Delegating permissions to your Amazon FSx service account](#).

Amazon FSx can't communicate with your Active Directory DNS servers or domain controllers

Joining an SVM to your self-managed Active Directory fails with the following error message:

Amazon FSx is unable to communicate with your Active Directory. This is because Amazon FSx can't reach the DNS servers provided or domain controllers for your domain. To fix this problem, update your storage virtual machine's Active Directory configuration with valid DNS servers and a networking configuration that allows traffic to flow from the storage virtual machine to the domain controller.

To resolve this issue, use the following procedure:

1. If only some of the domain controllers in your Active Directory are reachable, for example due to geographical limitations or firewalls, you can add preferred domain controllers. Using this option, Amazon FSx attempts to contact the preferred domain controllers. Add preferred domain controllers using the [vserver cifs domain preferred-dc add](#) NetApp ONTAP CLI command, as follows:

- a. To access the NetApp ONTAP CLI, establish an SSH session on the management port of the Amazon FSx for NetApp ONTAP file system by running the following command. Replace *management_endpoint_ip* with the IP address of the file system's management port.

```
[~]$ ssh fsxadmin@management_endpoint_ip
```

For more information, see [Managing file systems with the ONTAP CLI](#).

- b. Enter the following command, where:
 - `-vserver vserver_name` specifies the storage virtual machine (SVM) name.
 - `-domain domain_name` specifies the fully qualified Active Directory name (FQDN) of the domain to which the specified domain controllers belong.
 - `-preferred-dc IP_address,...` specifies one or more IP addresses of the preferred domain controllers, as a comma-delimited list, in order of preference.

```
FsxId123456789::> vserver cifs domain preferred-dc add -vserver vserver_name -  
domain domain_name -preferred-dc IP_address, ...+
```

The following command adds domain controllers 172.17.102.25 and 172.17.102.24 to the list of preferred domain controllers that the SMB server on SVM vs1 uses to manage external access to the cifs.lab.example.com domain.

```
FsxId123456789::> vserver cifs domain preferred-dc add -vserver vs1 -domain  
cifs.lab.example.com -preferred-dc 172.17.102.25,172.17.102.24
```

2. Check to see if your Domain Controller can be resolved with DNS. Use the [vserver services access-check dns forward-lookup](#) NetApp ONTAP CLI command to return

the IP address of a hostname based on the lookup on the DNS server specified or the vservers's DNS configuration.

- a. To access the NetApp ONTAP CLI, establish an SSH session on the management port of the Amazon FSx for NetApp ONTAP file system by running the following command. Replace *management_endpoint_ip* with the IP address of the file system's management port.

```
[~]$ ssh fsxadmin@management_endpoint_ip
```

For more information, see [Managing file systems with the ONTAP CLI](#).

- b. Enter the ONTAP CLI advanced mode using the following command.

```
FsxId123456789::> set adv
```

- c. Enter the following command, where:

- `-vserver vservice_name` specifies the storage virtual machine (SVM) name.
- `-hostname host_name` specifies the hostname to look up on the DNS server.
- `-node node_name` specifies the name of the node on which the command is executed.
- `-lookup-type` specifies the type of IP address to be looked up on the DNS server, default is all.

```
FsxId123456789::> vservice services access-check dns forward-lookup \  
-vserver vservice_name -node node_name \  
-domains domain_name -name-servers dns_server_ip_address \  
-hostname host_name
```

3. Review the [information you need to have](#) when joining an SVM to an AD.
4. Review the [networking requirements](#) when joining an SVM to an AD.
5. Use the procedure described in [Network configuration requirements](#) to update your SVM's AD configuration using the correct IP addresses for your AD DNS servers.

Amazon FSx can't communicate with your Active Directory because of a invalid Active Directory domain name.

Joining an SVM to your self-managed Active Directory fails with the following error message:

Amazon FSx has detected the provided FQDN is invalid. To fix this problem, update your storage virtual machine's Active Directory configuration with an FQDN that adheres to configuration requirements.

To resolve this issue, use the following procedure:

1. Review the on-premises Active Directory domain name requirements described in [Information needed when joining an SVM to an Active Directory](#). Make sure that the AD you are attempting to join meets that requirement.
2. Use the procedure described in [Joining an SVM to an Active Directory using the Amazon Web Services Management Console, Amazon CLI and API](#) and reattempt joining your SVM to an AD. Be sure to use the correct format for the AD domain's FQDN.

The service account can't access the administrators group specified in the SVM Active Directory configuration

Joining an SVM to your self-managed Active Directory fails with the following error message:

Amazon FSx is unable to apply your Active Directory configuration. This is because the administrators group you provided either doesn't exist or isn't accessible to the service account you provided. To fix this problem, ensure that your networking configuration allows traffic from the SVM to your Active Directory's domain controller(s) and DNS servers. Then update your SVM's Active Directory configuration, providing your Active Directory's DNS servers and, specifying an administrators group in the domain that is accessible to the service account provided.

To resolve this issue, do the following:

1. Review the information about [providing a domain group](#) to perform administrative actions on your SVM. Make sure that you are using the correct name of the AD domain administrators group.
2. Use the procedure described in [Joining an SVM to an Active Directory using the Amazon Web Services Management Console, Amazon CLI and API](#) and reattempt joining your SVM to an AD.

Amazon FSx can't connect to the Active Directory domain controllers because the organizational unit specified doesn't exist or isn't accessible

Joining an SVM to your self-managed Active Directory fails with the following error message:

Amazon FSx is unable to establish a connection with your Active Directory. This is because the organizational unit you specified either doesn't exist or isn't accessible to the service account provided. To fix this problem, update your storage virtual machine's Active Directory configuration, specifying an organizational unit to which the service account has permissions to join.

To resolve this issue, do the following:

1. Review the [prerequisites for joining an SVM to an AD](#).
2. Review the [information that you need to have](#) when joining an SVM to an AD.
3. Reattempt joining the SVM to the AD using [this procedure](#) with the correct organization unit.

You can't delete a storage virtual machine or volume

Each FSx for ONTAP file system can contain one or more storage virtual machines (SVMs), and each SVM can contain one or more volumes. When you delete a resource, you must first ensure that all its children have been deleted. For example, before deleting an SVM, you must first delete all the non-root volumes in the SVM.

Important

You can only delete storage virtual machines by using the Amazon FSx console, API, and CLI. You can only delete volumes using the Amazon FSx console, API, or CLI if the volume has Amazon FSx backups enabled.

To help protect your data and configuration, Amazon FSx prevents the deletion of SVMs and volumes in certain circumstances. If you attempt to delete an SVM or volume, and your deletion request doesn't succeed, Amazon FSx provides you with information in the Amazon console, Amazon Command Line Interface (Amazon CLI), and API regarding why the resource wasn't deleted. After you have addressed the cause of the deletion failure, you can retry the deletion request.


Topics

- [Identifying failed deletions](#)
- [SVM deletion: Route tables inaccessible](#)
- [SVM deletion: Peer relationship](#)
- [SVM or volume deletion: SnapMirror](#)
- [SVM deletion: Kerberos-enabled LIF](#)
- [SVM deletion: Other reason](#)
- [Volume deletion: FlexCache relationship](#)

Identifying failed deletions

When you delete an Amazon FSx SVM or volume, you typically see the resource's Lifecycle state transition to DELETING for up to a few minutes before the resource disappears from the Amazon FSx console, CLI, and API.

If you attempt to delete a resource and its Lifecycle state transitions from to DELETING and then back to CREATED, this behavior indicates that the resource didn't successfully delete. In this case, Amazon FSx reports an alert icon in the console next to the CREATED Lifecycle state. Choosing the alert icon displays the reason for the unsuccessful deletion, as shown in the following example.




The most common reasons why Amazon FSx prevents SVM and volume deletion are provided in the following sections, with step-by-step instructions on how to resolve these issues.

SVM deletion: Route tables inaccessible

Each FSx for ONTAP file system creates one or more route table entries to provide automatic failover and fail back across Availability Zones. By default, these route table entries are created in your VPC's default route table. You can optionally specify one or more non-default route tables where FSx for ONTAP interfaces can be created. Amazon FSx tags each route table that it associated with a file system with an AmazonFSx tag, and if this tag is removed, it can prevent Amazon FSx from being able to delete resources. If this situation occurs, you see the following LifecycleTransitionReason:

```
Amazon FSx is unable to complete the requested storage virtual machine operation because of an inability to access one or more of the route tables associated with your file system. Please contact Amazon Web Services Support.
```

You can find your file system's route tables in the Amazon FSx console by navigating to the file system's summary page, under the **Network & security** tab:



Choosing the route tables link takes you to your route tables. Next, verify that each of the route tables associated with your file system is tagged with this key-value pair:

Key: AmazonFSx
Value: ManagedByAmazonFSx

Tags	
<input type="text" value="Search tags"/>	
Key	Value
Name	Default
AmazonFSx	ManagedByAmazonFSx

If this tag isn't present, recreate it, and then try to delete the SVM again.

SVM deletion: Peer relationship

If you're attempting to delete an SVM or volume that's part of a peer relationship, you must first delete the peer relationship before you delete the SVM or volume. This requirement prevents the peered SVMs from becoming unhealthy. If your SVM can't be deleted because of a peer relationship, you see the following `LifecycleTransitionReason`:

Amazon FSx is unable to delete the storage virtual machine because it is part of a SVM peer or transition peer relationship. Please delete the relationship and retry.

You can delete SVM peer relationships through the ONTAP CLI. To access the ONTAP CLI, follow the steps in [Managing file systems with the ONTAP CLI](#). Using the ONTAP CLI, take the following steps.

1. Check for SVM peer relationships by using the following command. Replace *svm_name* with the name of your SVM.

```
FsxId123456789::> vserver peer show -vserver svm_name
```

If this command is successful, you'll see output similar to the following:

```

Vserver      Peer      Peer      Peering      Remote
Vserver      Vserver   State     Peer Cluster Applications Vserver
-----
svm_name    test2     peered    FsxId02d81fef0d84734b6
                                     snapmirror  fsxDest
svm_name    test3     peered    FsxId02d81fef0d84734b6
                                     snapmirror  fsxDest
2 entries were displayed.
```

2. Delete each SVM peer relationship by using the following command. Replace *svm_name*, and *remote_svm_name* with your actual values.

```
FsxId123456789abcdef::> vserver peer delete -vserver svm_name -peer-
vserver remote_svm_name
```

If this command is successful, you'll see the following output:

```
Info: 'vserver peer delete' command is successful.
```

SVM or volume deletion: SnapMirror

Just as you can't delete an SVM with a peer relationship without first deleting the peer relationship (see [SVM deletion: Peer relationship](#)), you can't delete an SVM that has a SnapMirror relationship without first deleting the SnapMirror relationship. To delete the SnapMirror relationship, use the ONTAP CLI to take the following steps on the file system that's the destination of the SnapMirror relationship. To access the ONTAP CLI, follow the steps in [Managing file systems with the ONTAP CLI](#).

Note

Amazon FSx backups use SnapMirror to create point-in-time, incremental backups of your file system's volumes. You can't delete this SnapMirror relationship for your backups in the ONTAP CLI. However, this relationship is automatically deleted when you delete a volume through the Amazon CLI, API, or console.

1. List your SnapMirror relationships on the destination file system by using the following command. Replace *svm_name* with the name of your SVM.

```
FsxId123456789abcdef::> snapmirror show -vserver svm_name
```

If this command is successful, you'll see output similar to the following:

Source Path	Type	Destination Path	Mirror State	Relationship Status	Total Progress	Last Healthy Updated
-----	----	-----	-----	-----	-----	-----

```
sourceSvm:sourceVol
      XDP  destSvm:destVol Snapmirrored
                               Idle           -           true    -
```

2. Delete your SnapMirror relationship by running the following command on the destination file system.

```
FsxId123456789abcdef::> snapmirror release -destination-path destSvm:destVol -
source-path sourceSvm:sourceVol -force true
```

SVM deletion: Kerberos-enabled LIF

If you are attempting to delete an SVM that has a logical interface (LIF) with Kerberos enabled, you must first disable Kerberos on that LIF before deleting the SVM.

You can disable Kerberos on a LIF through the ONTAP CLI. To access the ONTAP CLI, follow the steps in [Managing file systems with the ONTAP CLI](#).

1. Enter diagnostic mode in the ONTAP CLI by using the following command.

```
FsxId123456789abcdef::> set diag
```

When prompted to continue, enter **y**.

```
Warning: These diagnostic commands are for use by NetApp personnel only.
Do you want to continue? {y|n}: y
```

2. Check which interfaces have Kerberos enabled. Replace *svm_name* with the name of your SVM.

```
FsxId123456789abcdef::> kerberos interface show -vserver svm_name
```

If this command is successful, you'll see output similar to the following:

```
(vserver nfs kerberos interface show)
      Logical
Vserver  Interface      Address          Kerberos SPN
-----  -
svm_name  nfs_smb_management_1
                               10.19.153.48   enabled
```

```
5 entries were displayed.
```

3. Disable the Kerberos LIF by using the following command. Replace *svm_name* with the name of your SVM. You'll need to provide the Active Directory username and password that you used to join this SVM to your Active Directory.

```
FsxId123456789abcdef:> kerberos interface disable -vserver svm_name -lif
nfs_smb_management_1
```

If this command is successful, you'll see the following output. Provide the Active Directory username and password that you used to join this SVM to your Active Directory. When prompted to continue, enter *y*.

```
(vserver nfs kerberos interface disable)
Username: admin
Password: *****

Warning: This command deletes the service principal name from the machine account
on the KDC.
Do you want to continue? {y|n}: y

Disabled Kerberos on LIF "nfs_smb_management_1" in Vserver "svm_name".
```

4. Verify that Kerberos is disabled on the SVM by using the following command. Replace *svm_name* with the name of your SVM.

```
FsxId123456789abcdef:> kerberos interface show -vserver svm_name
```

If this command is successful, you'll see output similar to the following:

```
(vserver nfs kerberos interface show)
          Logical
Vserver   Interface      Address          Kerberos SPN
-----
svm_name  nfs_smb_management_1
                               10.19.153.48   disabled
5 entries were displayed.
```

5. If the interface is shown as disabled, try to delete the SVM again through the Amazon CLI, API, or console.

If you weren't able to delete the LIF by using the preceding commands, you can force-delete the Kerberos LIF by using the following command. Replace *svm_name* with the name of your SVM.

Important

The following command can strand the computer object of your SVM on your Active Directory.

```
FsxId123456789abcdef::> kerberos interface disable -vserver svm_name -lif  
nfs_smb_management_1 -force true
```

If this command is successful, you'll see output similar to the following. When prompted to continue, enter **y**.

```
(vserver nfs kerberos interface disable)  
  
Warning: Kerberos configuration for LIF "nfs_smb_management_1" in Vserver  
"svm_name" will be deleted.  
The corresponding account on the KDC will not be deleted. Do you want to continue?  
{y|n}: y
```

SVM deletion: Other reason

FSx for ONTAP SVMs create a computer object in your Active Directory when they join your Active Directory. In some cases, you may want to manually unjoin an SVM from your Active Directory by using the ONTAP CLI. To access the ONTAP CLI, follow the steps in [Managing file systems with the ONTAP CLI](#), logging into the ONTAP CLI at the file system level with `fsxadmin` credentials. Using the ONTAP CLI, take the following steps to unjoin an SVM from your Active Directory.

Important

This procedure can strand the computer object of your SVM on your Active Directory.

1. Enter advanced mode in the ONTAP CLI by using the following command.



```
FsxId123456789abcdef::> set adv
```

After running this command, you'll see this output. Enter **y** to continue.

```
Warning: These advanced commands are potentially dangerous; use them only when
        directed to do so by NetApp personnel.
        Do you want to continue? {y|n}: y
```

2. Delete the DNS for your Active Directory by using the following command. Replace *svm_name* with the name of your SVM.

```
FsxId123456789abcdef::> vserver services name-service dns dynamic-update record
delete -vserver svm_name -lif nfs_smb_management_1
```

 **Note**

If the DNS record has already been deleted or if the DNS server is unreachable, this command fails. If that happens, continue with the next step.

3. Disable the DNS by using the following command. Replace *svm_name* with the name of your SVM.

```
FsxId123456789abcdef::> vserver services name-service dns dynamic-update modify -
vserver svm_name -is-enabled false -use-secure false
```

If this command is successful, you'll see the following output:

```
Warning: DNS updates for Vserver "svm_name" are now disabled.
        Any LIFs that are subsequently modified or deleted
        can result in a stale DNS entry on the DNS server,
        even when DNS updates are enabled again.
```

4. Unjoin the device from Active Directory. Replace *svm_name* with the name of your SVM.

```
FsxId123456789abcdef::> vserver cifs delete -vserver svm_name
```

After running this command, you'll see the following output, where *CORP.EXAMPLE.COM* is replaced with the name of your domain. When prompted, enter your user name and password. When asked if you want to delete the server, enter *y*.

```
In order to delete an Active Directory machine account for the CIFS server,
you must supply the name and password of a Windows account with sufficient
privileges to remove computers from the "CORP.EXAMPLE.COM" domain.
Enter the user name: admin
Enter the password:
Warning: There are one or more shares associated with this CIFS server
    Do you really want to delete this CIFS server and all its shares? {y|n}: y
Warning: Unable to delete the Active Directory computer account for this CIFS
server.
    Do you want to continue with CIFS server deletion anyway? {y|n}: y
```

Volume deletion: FlexCache relationship

You can't delete volumes that are the origin volumes for a FlexCache relationship unless you delete the cache relationship first. To determine which volumes have a FlexCache relationship, you can use the ONTAP CLI. To access the ONTAP CLI, follow the steps in [Managing file systems with the ONTAP CLI](#).

1. Check for FlexCache relationships by using the following command.

```
FsxId123456789abcdef::> volume flexcache origin show-caches
```

2. Delete any cache relationships by using the following command. Replace *dest_svm_name*, and *dest_vol_name* with your actual values.

```
FsxId123456789abcdef::> volume flexcache delete -server dest_svm_name -
volume dest_vol_name
```

3. After you've deleted the cache relationship, try to delete your SVM through the Amazon CLI, API, or console again.

Automatic daily backups fail due to insufficient volume capacity

Automatic daily backups of your volume fails with the following message:

```
Amazon FSx could not create a backup of your volume because the backup snapshot was deleted.
```

Automatic daily backups are failing because there is insufficient free storage capacity on the volume. To mitigate this condition, you will need to free up storage capacity on the volume. You can accomplish this using one or more of the following options, depending on your situation:

- [Increase the volume's storage capacity](#)
- [Increase the volume's snapshot reserve](#)
- [Disable snapshot auto-delete](#)
- Don't delete the backup-snapshot using the ONTAP CLI

You have insufficient volume capacity

If you are running out of space on your volumes, you can use the procedures shown here to diagnose and resolve the situation.

Topics

- [Determine how your volume storage capacity is being used](#)
- [Increasing a volume's storage capacity](#)
- [Using volume autosizing](#)
- [Your file system's primary storage is full](#)
- [Deleting snapshots](#)
- [Increasing a volume's maximum file capacity](#)

Determine how your volume storage capacity is being used

You can see how your volume's storage capacity is being consumed by using the `volume show-space` NetApp ONTAP CLI command. This information can help you make decisions about how

to reclaim or conserve volume storage capacity. For more information, see [To monitor a volume's storage capacity \(console\)](#).

Increasing a volume's storage capacity

You can increase a volume's storage capacity by using the Amazon FSx console, Amazon CLI, and Amazon FSx API. For more information about updating a volume with an increased capacity, see [Updating a volume](#).

Alternatively, you can increase a volume's storage capacity using the [volume modify](#) NetApp ONTAP CLI command. For more information, see [To change a volume's storage capacity \(console\)](#).

Using volume autosizing

You can use volume autosizing so that a volume automatically grows by a specified amount, or to a specified size when it reaches a used space threshold. You can do this for FlexVol volume types, which is the default volume type for FSx for ONTAP, using the [volume autosize](#) NetApp ONTAP CLI command. For more information, see [Enabling volume autosizing](#).

Your file system's primary storage is full

If your FSx for ONTAP file system's primary storage is full, you cannot add any more data to the volumes in your file system, even if a volume is showing that it has enough available storage capacity. You can view the amount of available primary storage capacity in the **Monitoring & performance** tab on the file system details page in the Amazon FSx console. For more information, see [Monitoring SSD storage utilization](#)

To resolve this issue, you can increase the size of your file system's primary storage tier. For more information, see [Updating file system SSD storage and IOPS](#).

Deleting snapshots

Snapshots are enabled by default on your volumes, using the default snapshot policy. Snapshots are stored in the `.snapshot` directory at the root of a volume. You can manage volume storage capacity with respect to snapshots in the following ways:

- [Manually delete snapshots](#) – reclaim storage capacity by deleting snapshots manually.
- [Create a snapshot autodelete policy](#) – create a policy that deletes snapshots more aggressively than the default snapshot policy.

- [Turn off automatic snapshots](#) – conserve storage capacity by turning off automatic snapshots.

For more information about deleting snapshots and managing snapshot policies to conserve storage capacity, see [Deleting snapshots](#).

Increasing a volume's maximum file capacity

An FSx for ONTAP volume can run out of file capacity when the number of available inodes, or file pointers, is exhausted. By default, the number of available inodes on a volume is 1 for every 32KiB of volume size. For more information, see [Volume file capacity](#).

The number of inodes in a volume increases commensurately with the volume's storage capacity, up to a threshold of 648 GiB. By default, volumes that have storage capacity of 648 GiB or more all have the same number of inodes, 21,251,126. To view a volume's maximum file capacity, see [Viewing a volume's file capacity](#).

If you create a volume larger than 648 GiB, and you want to have more than 21,251,126 inodes, you must increase the maximum number of files on the volume manually. If your volume is running out of storage capacity, you can check its maximum file capacity. If it's nearing its file capacity, you can manually increase it. For more information, see [To increase the maximum number of files on a volume \(ONTAP CLI\)](#).

Troubleshooting network issues

If you are experiencing network issues, you can use the procedures shown here to diagnose the problem.

You want to capture a packet trace

Packet tracing is the process of verifying the path of a packet through the layers to its destination. You control the packet tracing process with the following NetApp ONTAP CLI commands:

- **network tcpdump start** – Starts packet tracing
- **network tcpdump show** – Shows currently running packet traces
- **network tcpdump stop** – Stops a running packet trace

These commands are available to users who have the `fsxadmin` role on your file system.

To capture a packet trace from your file system

1. To SSH into the NetApp ONTAP CLI of your file system, follow the steps documented in the [Using the NetApp ONTAP CLI](#) section of the *Amazon FSx for NetApp ONTAP User Guide*.

```
ssh fsxadmin@file-system-management-endpoint-ip-address
```

2. Enter the diagnostic privilege level in the ONTAP CLI by using the following command.

```
::> set diag
```

When prompted to continue, enter y.

```
Warning: These diagnostic commands are for use by NetApp personnel only.
Do you want to continue? {y|n}: y
```

3. Identify the location on your file system where you want to save your packet trace. The volume must be online and must be mounted in the namespace with a valid junction path. Use the following command to check for volumes that fulfill those criteria:

```
::*> volume show -junction-path !- -fields junction-path
vserver volume    junction-path
-----
fsx      test_vol1 /test_vol1
fsx      test_vol2 /test_vol2
fsx      test_vol2 /test_vol3
```

4. Start the trace with the minimum required arguments. Replace the following:
 - Replace *node_name* with the name of the node (for example, FsxId01234567890abcdef-01).
 - Replace *svm_name* with the name of your storage virtual machine (for example, fsx).
 - Replace *junction_path_name* with the name of the volume (for example, test-vol1).

```
::*> debug network tcpdump start -node node_name -ip-space Default -pass-through "-i
e0e -w /clus/svm_name/junction_path_name"
Info: Started network trace on interface "e0e"
Warning: Snapshots should be disabled on the tcpdump destination volume while
packet traces are occurring. Use the
```

```
"volume modify -snapshot-policy none -vserver fsx -volume test_vol1" command to
disable Snapshots on the
tcpdump destination volume.
```

Important

Packet traces can only be captured on the e0e interface and in the Default IP space. In FSx for ONTAP, all network traffic uses the e0e interface.

When using packet tracing, keep the following in mind:

- When starting a packet trace, you must include the path to where you want to store the trace files, in this format: `/clus/svm_name/junction-path-name`
- Optionally, provide the file name for the packet trace. If the `filter_name` is not specified, it is automatically generated in the form: `node-name_port-name_yyyymmdd_hhmmss.trc`
- If rolling traces are specified, the `filter_name` is suffixed with a number that indicates the position in the rotation sequence.
- The ONTAP CLI also accepts the following optional **-pass-through** arguments:

```
-B, --buffer-size=<KiB>
-c <number_of_packets>
-C <file_size-mB>
-F <filter_expression_filename>
-G <rotate_seconds>
--time-stamp-precision {micro|nano}
-Q, --direction {in|out|inout}
-s, --snapshot-length=<bytes>
-U, --packet-buffered
-W <rotate_file_count>
<filter-expression>
```

- For information about filter expressions, see [pcap-filter\(7\) man page](#).

5. View the traces in progress:

```
::*> debug network tcpdump show
Node           IPspace  Port    Filename
-----
```

```
FsxId123456789abcdef-01 Default e0e /clus/fsx/test_vol1/  
FsxId123456789abcdef-01_e0e_20230605_181451.trc
```

6. Stop the trace:

```
::*> debug network tcpdump stop -node FsxId123456789abcdef-01 -ipspace Default -  
port e0e  
Info: Stopped network trace on interface "e0e"
```

7. Return to the admin privilege level:

```
::*> set -priv admin  
::>
```

8. Access the packet traces.

Your packet traces are stored in the volume that you specified using the **debug network tcpdump start** command, and can be accessed via the NFS export or an SMB share that corresponds with that volume.

For more information about capturing packet traces, see [How to use debug network tcpdump in ONTAP 9.10+](#) in the NetApp Knowledge Base.

Document History for Amazon FSx for NetApp ONTAP

- **API version:** 2018-03-01
- **Latest documentation update:** March 4, 2024

The following table describes important changes to the *Amazon FSx NetApp ONTAP User Guide*. For notifications about documentation updates, you can subscribe to the RSS feed.

Change	Description	Date
Support added for 12 HA pairs in scale-out file systems	Amazon FSx for NetApp ONTAP added support for 12 HA pairs in scale-out file systems. File systems with 12 HA pairs can deliver up to 72 GBps of throughput capacity and 2,400,000 SSD IOPS across 12 high-availability (HA) pairs. For more information, see High-availability (HA) pairs and Amazon FSx for NetApp ONTAP performance .	March 4, 2024
Support added for cloud write mode	Amazon FSx for NetApp ONTAP added support for cloud write mode for volumes. For more information, see Enabling cloud write mode on a volume .	February 6, 2024
Support added for backing up FlexGroup volumes with Amazon Backup	You can now use Amazon Backup to back up and restore FlexGroup volumes on your FSx for ONTAP file systems. For more information, see	January 11, 2024

[Using Amazon Backup with Amazon FSx.](#)

[Amazon FSx updated the AmazonFSxFullAccess, AmazonFSxConsoleFullAccess, AmazonFSxReadOnlyAccess, AmazonFSxConsoleReadOnlyAccess, and AmazonFSxServiceRolePolicy Amazon managed policies](#)

Amazon FSx updated the AmazonFSxFullAccess, AmazonFSxConsoleFullAccess, AmazonFSxReadOnlyAccess, AmazonFSxConsoleReadOnlyAccess, and AmazonFSxServiceRolePolicy policies to add the `ec2:GetSecurityGroupsForVpc` permission. For more information, see [Amazon FSx updates to Amazon managed policies](#).

January 9, 2024

[Amazon FSx updated the AmazonFSxFullAccess and the AmazonFSxConsoleFullAccess Amazon managed policies](#)

Amazon FSx updated the AmazonFSxFullAccess and AmazonFSxConsoleFullAccess policies to add the `ManageCrossAccountDataReplication` action. For more information, see [Amazon FSx updates to Amazon managed policies](#).

December 20, 2023

[Support added for scale-out metrics](#)

FSx for ONTAP now provides Amazon CloudWatch metrics for file systems with multiple HA pairs. For more information, see [Scale-out file system metrics](#).

November 26, 2023

[Support added for scale-out file systems](#)

Amazon FSx for NetApp ONTAP added support for scale-out file systems that can deliver up to 36 GBps of throughput capacity and 1,200,000 SSD IOPS across six high-availability (HA) pairs. For more information, see [High-availability \(HA\) pairs](#) and [Amazon FSx for NetApp ONTAP performance](#).

November 26, 2023

[Support added for FlexGroup volumes](#)

Amazon FSx for NetApp ONTAP added support for FlexGroup volumes. For more information, see [Volume styles](#).

November 26, 2023

[Shared VPC support added for Multi-AZ file systems](#)

Participant accounts can now create Multi-AZ file systems in a VPC that has been shared with them. Owner accounts can manage this feature in the Amazon FSx console, CLI, and API. For more information, see [Creating FSx for ONTAP file systems in shared subnets](#)

November 26, 2023

[Amazon FSx updated the AmazonFSxFullAccess and the AmazonFSxConsoleFullAccess Amazon managed policies](#)

Amazon FSx updated the AmazonFSxFullAccess and AmazonFSxConsoleFullAccess policies to add the fsx:CopySnapshotAndUpdateVolume permission. For more information, see [Amazon FSx updates to Amazon managed policies](#).

November 26, 2023

[Amazon FSx updated the AmazonFSxFullAccess and the AmazonFSxConsoleFullAccess Amazon managed policies](#)

Amazon FSx updated the AmazonFSxFullAccess and AmazonFSxConsoleFullAccess policies to add the fsx:DescribeSharedVPCConfiguration and fsx:UpdateSharedVPCConfiguration permissions. For more information, see [Amazon FSx updates to Amazon managed policies](#).

November 14, 2023

[Support added for creating additional ONTAP roles and users](#)

Amazon FSx for NetApp ONTAP now supports creating additional ONTAP roles and users to define user capabilities and privileges when using the ONTAP CLI and REST API. For more information, see [Roles and users in Amazon FSx for NetApp ONTAP](#).

September 6, 2023

Support added for additional CloudWatch metrics and an enhanced monitoring dashboard	FSx for ONTAP now provides additional performance metrics and an enhanced monitoring dashboard for improved visibility into file system activity. For more information, see Monitoring with CloudWatch .	August 17, 2023
Amazon FSx updated the AmazonFSxServiceRolePolicy Amazon managed policy	Amazon FSx updated the <code>cloudwatch:PutMetricData</code> permission in the <code>AmazonFSxServiceRolePolicy</code> . For more information, see Amazon FSx updates to Amazon managed policies .	July 24, 2023
Support added for using NetApp System Manager directly	You can manage your FSx for ONTAP file systems using System Manager directly from NetApp BlueXP. For more information, see Using NetApp System Manager with BlueXP .	July 13, 2023
Support added for monitoring EMS events	You can monitor FSx for ONTAP file system events using NetAPP ONTAP's native Events Management System (EMS). You can view EMS events using the NetApp ONTAP CLI. For more information, see Monitoring FSx for ONTAP EMS events .	July 13, 2023

[Support added for SnapLock](#)

FSx for ONTAP now supports SnapLock volumes. SnapLock allows you to protect your files by transitioning them to a write once, read many (WORM) state, which prevents modification or deletion for a specified retention period. FSx for ONTAP supports the Compliance and Enterprise retention modes with SnapLock. For more information, see [Working with SnapLock](#).

July 13, 2023

[Support added for IPsec encryption of data in transit](#)

FSx for ONTAP now supports using IPsec encryption to encrypt data in transit between file systems and connected clients. For more information, see [Configuring IPsec using PSK authentication](#) and [Configuring IPsec using certificate authentication](#).

July 13, 2023

[Maximum volume size has increased](#)

FSx for ONTAP updated the maximum size of a volume from 100 TB to 300 TB. For more information, see [Turn on volume autosizing](#).

July 13, 2023

Amazon FSx updated the AmazonFSxFullAccess Amazon managed policy	Amazon FSx updated the AmazonFSxFullAccess policy to remove the fsx:* permission and add specific fsx actions. For more information, see AmazonFSx FullAccess policy.	July 13, 2023
Amazon FSx updated the AmazonFSxConsoleFullAccess Amazon managed policy	Amazon FSx updated the AmazonFSxConsoleFullAccess policy to remove the fsx:* permission and add specific fsx actions. For more information, see AmazonFSx ConsoleFullAccess policy.	July 13, 2023
Support added for joining existing storage virtual machines to an Active Directory	You can join existing storage virtual machines to an Active Directory using the Amazon Web Services Management Console, Amazon CLI and API. For more information, see Joining an SVM to an Active Directory .	June 13, 2023
Support for NVMe read cache added for Single-AZ file systems	NVMe read cache is now supported for Single-AZ file systems created after November 28, 2022 with at least 2 GBps of throughput capacity in US East (Ohio) Region, US East (N. Virginia) Region, US West (Oregon) Region, and Europe (Ireland) . For more information, see Impact of deployment type on performance .	November 28, 2022

[Support added for using in-VPC IP address ranges to create Multi-AZ file systems](#)

You can now create Multi-AZ FSx for ONTAP file systems by specifying endpoints that are within the IP address range of your VPC. For more information, see [Creating FSx for ONTAP file systems](#).

November 28, 2022

[Support added for updating VPC route tables on Multi-AZ file systems](#)

You can now associate (add) a new VPC route table to an existing Multi-AZ FSx for ONTAP file system or disassociate (remove) an existing VPC route table from an existing Multi-AZ FSx for ONTAP file system. For more information, see [Updating a file system](#).

November 28, 2022

[Support added for encryption of data in transit with Amazon Nitro System](#)

Data in transit is encrypted automatically when accessed from supported Amazon EC2 instances in US East (Ohio) Region, US East (N. Virginia) Region, US West (Oregon) Region, and Europe (Ireland) . For more information, see [Encrypting data in transit with Amazon Nitro System](#).

November 28, 2022

[Support added for creating DP volumes](#)

You can now create DP (data-protection) volumes by using the Amazon FSx console, Amazon CLI, or Amazon FSx API. You can use DP volumes as the destination of a NetApp SnapMirror or SnapVault relationship, when you want to migrate or protect a single volume's data. For more information, see [Volume types](#).

November 28, 2022

[Support added for copying volume tags to backups](#)

You can now enable CopyTagsToBackups in the Amazon CLI or Amazon FSx API to automatically copy tags from your volumes to backups. For more information, see [Copying tags to backups](#).

November 28, 2022

[Support added for choosing a snapshot policy](#)

You can now choose from three built-in snapshot policies when creating or updating a volume using the Amazon FSx console, Amazon CLI, or Amazon FSx API. You can also select a custom snapshot policy that you made in the ONTAP CLI or REST API. For more information, see [Snapshot policies](#).

November 28, 2022

[Support added for additional file system throughput capacity option](#)

FSx for ONTAP now supports 4,096 MBps of throughput capacity for file systems created after November 28, 2022 in US East (Ohio) Region, US East (N. Virginia) Region, US West (Oregon) Region, and Europe (Ireland). For more information, see [Impact of throughput capacity on performance](#).

November 28, 2022

[Support added for additional SSD IOPS](#)

FSx for ONTAP now supports 160,000 SSD IOPS for file systems created after November 28, 2022 in US East (Ohio) Region, US East (N. Virginia) Region, US West (Oregon) Region, and Europe (Ireland). For more information, see [Impact of throughput capacity on performance](#).

November 28, 2022

[Support added for using FSx for ONTAP as an external datastore for VMware Cloud on Amazon](#)

You can use FSx for ONTAP as an external datastore for VMware Cloud on Amazon Software-Defined Data Centers (SDDCs). This added support provides flexibility to scale storage up or down independently from compute resources for VMware Cloud on Amazon workloads. For more information, see [Using VMware Cloud with FSx for ONTAP](#).

August 30, 2022

[Automatically increase a file system's storage capacity](#)

Use an Amazon-developed customizable Amazon CloudFormation template to automatically increase your file system's storage capacity when the amount of used SSD storage capacity exceeds a threshold that you specify. For more information, see [Increasing SSD storage capacity dynamically](#).

June 3, 2022

[Amazon FSx is now integrated with Amazon Backup](#)

You can now use Amazon Backup to back up and restore your FSx file systems in addition to using the native Amazon FSx backups. For more information, see [Using Amazon Backup with Amazon FSx](#).

May 18, 2022

[Support added for single Availability Zone ONTAP file system deployments](#)

You can create Single-AZ FSx for ONTAP file systems, which are designed to provide high availability and durability within a single Availability Zone (AZ). For more information, see [Choosing file system deployment](#).

April 13, 2022

[Support added for Amazon PrivateLink interface VPC endpoints](#)

You can now use interface VPC endpoints to access the Amazon FSx API from your VPC without sending traffic over the internet. For more information, see [Amazon FSx and interface VPC endpoints](#).

April 5, 2022

Support added for modifying throughput capacity for existing ONTAP file systems	You can now modify the throughput capacity that is available to your existing ONTAP file systems. For more information, see Managing throughput capacity .	March 30, 2022
Support added for SSD storage capacity and provisioned IOPS scaling	You can now increase the SSD storage capacity and provisioned IOPS for existing FSx for ONTAP file systems as your storage and IOPS requirements evolve. For more information, see Managing storage capacity and provisioned IOPS .	January 25, 2022
Support added for Amazon CloudWatch metrics	You can monitor your file system using Amazon CloudWatch, which collects and processes raw data from FSx for ONTAP into readable, near real-time metrics. For more information, see Monitoring with Amazon CloudWatch .	January 19, 2022
Support added for additional file system throughput options	FSx for ONTAP now supports 128 MB/s and 256 MB/s options for file system throughput. For more information, see Impact of throughput capacity on performance .	November 30, 2021

[Amazon FSx for NetApp ONTAP is now generally available](#)

FSx for ONTAP is a fully managed service that provides highly reliable, scalable, performant, and feature-rich file storage built on NetApp's ONTAP file system. It provides the familiar features, performance, capabilities, and APIs of NetApp file systems with the agility, scalability, and simplicity of a fully managed Amazon service.

September 2, 2021