Amazon FSx File Gateway User Guide

AWS Storage Gateway

API Version 2021-03-31

AWS Storage Gateway: Amazon FSx File Gateway User Guide

Table of Contents

Amazon FSx File Gateway is no longer available to new customers. Existing customers of FSx File Gateway can continue to use the service normally. For capabilities similar to FSx File Gateway, visit this blog post.

What is Amazon FSx File Gateway

Amazon FSx File Gateway (FSx File Gateway) is a new File Gateway type that provides low latency and efficient access to in-cloud FSx for Windows File Server file shares from your on-premises facility. If you maintain on-premises file storage because of latency or bandwidth requirements, you can instead use FSx File Gateway for seamless access to fully managed, highly reliable, and virtually unlimited Windows file shares provided in the AWS Cloud by FSx for Windows File Server.

Benefits of using Amazon FSx File Gateway

FSx File Gateway provides the following benefits:

- Helps eliminate on-premises file servers and consolidates all their data in AWS to take advantage of the scale and economics of cloud storage.
- Provides options that you can use for all your file workloads, including those that require on-premises access to cloud data.
- Applications that need to stay on premises can now experience the same low latency and high performance that they have in AWS, without taxing your networks or impacting the latencies experienced by your most demanding applications.

How Amazon FSx File Gateway works

To use Amazon FSx File Gateway (FSx File Gateway), you must have at least one Amazon FSx for Windows File Server file system. You must also have on-premises access to FSx for Windows File Server, either through a VPN or through an AWS Direct Connect connection. For more information about using Amazon FSx file systems, see What is Amazon FSx for Windows File Server?

You deploy the gateway into your on-premises environment as a virtual machine (VM) running on VMware ESXi, Microsoft Hyper-V, or Linux Kernel-based Virtual Machine (KVM), or as a hardware appliance that you order from your preferred reseller. You can also deploy the Storage Gateway VM in VMware Cloud on AWS, or as an AMI in Amazon EC2. After deploying your appliance, you activate the FSx File Gateway from the Storage Gateway console or through the Storage Gateway API.

After the Amazon FSx File Gateway is activated and can access FSx for Windows File Server, use the Storage Gateway console to join it to your Microsoft Active Directory domain. After the

gateway successfully joins a domain, you use the Storage Gateway console to attach the gateway to an existing FSx for Windows File Server. FSx for Windows File Server makes all the shares on the server available as shares on your Amazon FSx File Gateway. You can then use a client to browse and connect to the file shares on FSx File Gateway that correspond to the selected FSx File Gateway.

When the file shares are connected, you can read and write your files locally, while benefiting from all the features available on FSx for Windows File Server. FSx File Gateway maps local file shares and their contents to file shares stored remotely in FSx for Windows File Server. There is a 1:1 correspondence between the remote and locally visible files and their shares.

The following diagram provides an overview of file storage deployment for Storage Gateway.

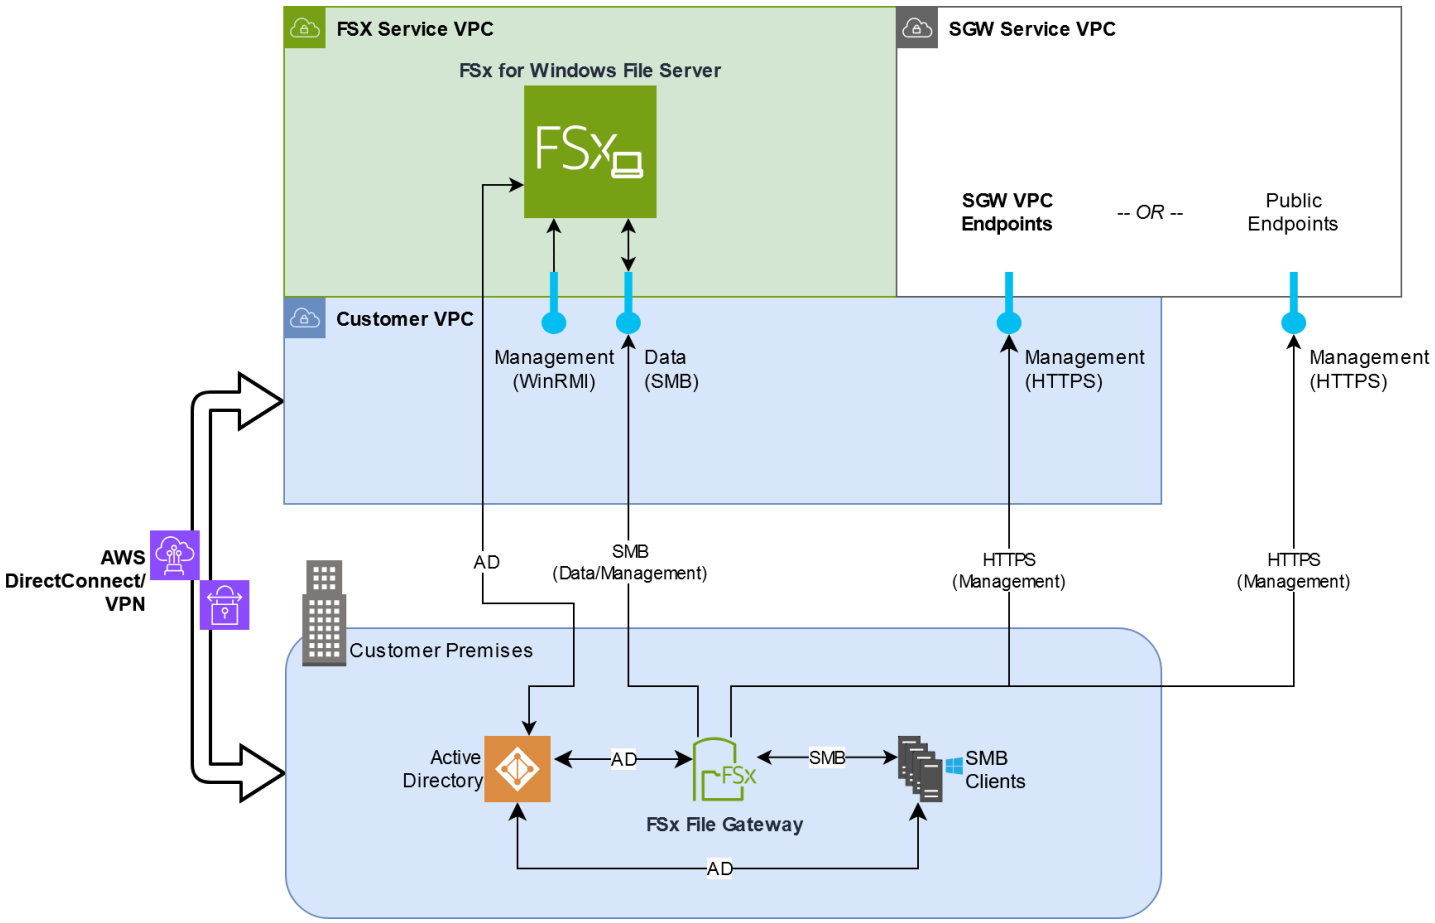

Note the following in the diagram:

- **AWS Direct Connect or a VPN** is needed to allow the FSx File Gateway to access the Amazon FSx file share using SMB and to allow the FSx for Windows File Server to join your on-premises Active Directory domain.

- **Amazon Virtual Private Cloud (Amazon VPC)** is needed to connect to the FSx for Windows File Server service VPC and the Storage Gateway service VPC using private endpoints. The FSx File Gateway can also connect to the public endpoints.

You can use Amazon FSx File Gateway in all AWS Regions where FSx for Windows File Server is available.

Getting Started with AWS Storage Gateway

This section provides instructions for getting started with AWS. You need an AWS account before you can start using AWS Storage Gateway. You can use an existing AWS account, or sign up for a new account. You also need an IAM user in your AWS account that belongs to a group with the necessary administrative permissions to perform Storage Gateway tasks. Users with the appropriate privileges can access the Storage Gateway console and Storage Gateway API to perform gateway deployment, configuration, and maintenance tasks. If you are a first-time user, we recommend that you review the Supported AWS regions and File Gateway setup requirements sections before you being working with Storage Gateway.

This section contains the following topics, which provide additional information about getting started with AWS Storage Gateway:

Topics

- Sign up for Amazon Web Services - Learn how to sign up for AWS and create an AWS account.
- Create an IAM user with administrator privileges - Learn how to create an IAM user with administrative privileges for your AWS account.
- Accessing AWS Storage Gateway - Learn how to access AWS Storage Gateway through the Storage Gateway console or programmatically using the AWS SDKs.
- AWS Regions that support Storage Gateway - Learn which AWS Regions you can use to store your data when you activate your gateway in Storage Gateway.

Sign up for Amazon Web Services

An AWS account is a fundamental requirement for accessing AWS services. Your AWS account is the basic container for all of the AWS resources you create as an AWS user. Your AWS account is also the basic security boundary for your AWS resources. Any resources that you create in your account are available to users who have credentials for the account. Before you can start using AWS Storage Gateway, you need to sign up for an AWS account.

If you do not have an AWS account, complete the following steps to create one.

To sign up for an AWS account

1. Open https://portal.aws.amazon.com/billing/signup.

2. Follow the online instructions.

   Part of the sign-up procedure involves receiving a phone call or text message and entering a verification code on the phone keypad.

   When you sign up for an AWS account, an *AWS account root user* is created. The root user has access to all AWS services and resources in the account. As a security best practice, assign administrative access to a user, and use only the root user to perform tasks that require root user access.

We also recommend that you require your users to use temporary credentials when accessing AWS. To provide temporary credentials, you can use federation and an identity provider, such as AWS IAM Identity Center. If your company already uses an identity provider, you can use it with federation to simplify how you provide access to the resources in your AWS account.

Create an IAM user with administrator privileges

After you create your AWS account, use the following steps to create an AWS Identity and Access Management (IAM) user for yourself, and then add that user to a group that has administrative permissions. For more information about using the AWS Identity and Access Management service to control access to Storage Gateway resources, see Identity and access management for AWS Storage Gateway.

To create an administrator user, choose one of the following options.

Choose one way to manage your administrator	To	By	You can also
In IAM Identity Center	Use short-term credentials to access AWS.	Following the instructions in Getting started in the *AWS IAM Identity Center User Guide*.	Configure programmatic access by Configuring the AWS CLI to use AWS IAM Identity Center in the *AWS*

Choose one way to manage your administrator	To	By	You can also
(Recommended)	This aligns with the security best practices. For information about best practices, see Security best practices in IAM in the *IAM User Guide*.		*Command Line Interface User Guide*.
In IAM (Not recommended)	Use long-term credentials to access AWS.	Following the instructions in Create an IAM user for emergency access in the *IAM User Guide*.	Configure programmatic access by Manage access keys for IAM users in the *IAM User Guide*.

Warning

IAM users have long-term credentials which present a security risk. To help mitigate this risk, we recommend that you provide these users with only the permissions they require to perform the task and that you remove these users when they are no longer needed.

Accessing AWS Storage Gateway

You can use the AWS Storage Gateway console to perform various gateway configuration and maintenance tasks, including activating or removing Storage Gateway hardware appliances from your deployment, creating, managing, and deleting the different types of gateways, attaching, managing, and detaching file systems, and monitoring the health and status of various elements of the Storage Gateway service. For simplicity and ease of use, this guide focuses on performing tasks

using the Storage Gateway console web interface. You can access the Storage Gateway console through your web browser at: https://console.aws.amazon.com/storagegateway/home/.

If you prefer a programmatic approach, you can use the AWS Storage Gateway Application Programming Interface (API) or Command Line Interface (CLI) to set up and manage the resources in your Storage Gateway deployment. For more information about actions, data types, and required syntax for the Storage Gateway API, see the Storage Gateway API Reference. For more information about the Storage Gateway CLI, see the AWS CLI Command Reference.

You can also use the AWS SDKs to develop applications that interact with Storage Gateway. The AWS SDKs for Java, .NET, and PHP wrap the underlying Storage Gateway API to simplify your programming tasks. For information about downloading the SDK libraries, see the AWS Developer Center.

For information about pricing, see AWS Storage Gateway pricing.

AWS Regions that support Storage Gateway

An AWS Region is a physical location in the world where AWS has multiple Availability Zones. Availability Zones consist of one or more discrete AWS data centers, each with redundant power, networking, and connectivity, housed in separate facilities. This means that each AWS Region is physically isolated and independent of the other Regions. Regions provide fault tolerance, stability, and resilience, and can also reduce latency. The resources that you create in one Region do not exist in any other Region unless you explicitly use a replication feature offered by an AWS service. For example, Amazon S3 and Amazon EC2 support cross-Region replication. Some services, such as AWS Identity and Access Management, do not have Regional resources. You can launch AWS resources in locations that meet your business requirements. For example, you might want to launch Amazon EC2 instances to host your AWS Storage Gateway appliances in an AWS Region in Europe to be closer to your European users, or to meet legal requirements. Your AWS account determines which of the Regions supported by a specific service are available for you to use.

Amazon FSx File Gateway stores file data in the AWS Region where your Amazon FSx file system is located. Before you start deploying your gateway, choose a Region in the upper-right corner of the Storage Gateway console.

- Amazon FSx File Gateway — For supported AWS Regions and a list of AWS service endpoints that you can use with Amazon FSx File Gateway, see Amazon FSx File Gateway endpoints and quotas in the *AWS General Reference*.

- Storage Gateway — For supported AWS Regions and a list of AWS service endpoints that you can use with Storage Gateway, see AWS Storage Gateway endpoints and quotas in the *AWS General Reference*.
- Storage Gateway Hardware Appliance — For supported Regions that you can use with the hardware appliance, see AWS Storage Gateway Hardware Appliance Regions in the *AWS General Reference*.

File Gateway setup requirements

Unless otherwise noted, the following requirements are common to all File Gateway types in AWS Storage Gateway. Your setup must meet the requirements in this section. Review the requirements that apply to your gateway setup before you deploy your gateway.

Topics

Prerequisites

Before you set up your Amazon FSx File Gateway (FSx File Gateway), you must meet the following prerequisites:

- Create and configure an FSx for Windows File Server file system. For instructions, see Step 1: Create Your File System in the *Amazon FSx for Windows File Server User Guide*.
- Configure Microsoft Active Directory (AD) and create an Active Directory service account with the requisite permissions. For more information, see Active Directory service account permission requirements.
- Ensure that there is sufficient network bandwidth between the gateway and AWS. A minimum of 100 Mbps is required to successfully download, activate, and update the gateway.
- Configure the connection you want to use for network traffic between AWS and the on-premises environment where you are deploying your gateway. You can connect using the public internet, private networking, a VPN, or AWS Direct Connect. If you want your gateway to communicate AWS through a private connection to an Amazon Virtual Private Cloud, set up the Amazon VPC before you set up your gateway.

- Make sure your gateway can resolve the name of your Active Directory Domain Controller. You can use DHCP in your Active Directory domain to handle resolution, or specify a DNS server manually from the Network Configuration settings menu in the gateway local console.

Hardware and storage requirements

The following sections provide information about the minimum required hardware and storage configurations for your gateway, and the minimum amount of disk space to allocate for the required storage.

Hardware requirements for on-premises VMs

When deploying your gateway on-premises, ensure that the underlying hardware on which you deploy the gateway virtual machine (VM) can dedicate the following minimum resources:

- Four virtual processors assigned to the VM
- 16 GiB of reserved RAM for File Gateways
- 80 GiB of disk space for installation of VM image and system data

Requirements for Amazon EC2 instance types

When deploying your gateway on Amazon Elastic Compute Cloud (Amazon EC2), the instance size must be at least **`xlarge`** for your gateway to function. However, for the compute-optimized instance family the size must be at least **`2xlarge`**.

> **Note**
>
> The Storage Gateway AMI is only compatible with x86-based instances that use Intel or AMD processors. ARM-based instances that use Graviton processors are not supported.

Use one of the following instance types recommended for your gateway type.

Recommended for File Gateway types

- General-purpose instance family – **m4, m5, m6, or m7** instance type. Choose the **xlarge** instance size or higher to meet the Storage Gateway processor and RAM requirements.

- Compute-optimized instance family – **c4, c5, c6, or c7** instance types. Choose the **2xlarge** instance size or higher to meet the Storage Gateway processor and RAM requirements.
- Memory-optimized instance family – **r3, r5, r6, or r7** instance types. Choose the **xlarge** instance size or higher to meet the Storage Gateway processor and RAM requirements.
- Storage-optimized instance family – **i3, i4 or i7** instance types. Choose the **xlarge** instance size or higher to meet the Storage Gateway processor and RAM requirements.

> **Note**
>
> When you launch your gateway in Amazon EC2 and the instance type you choose supports ephemeral storage, the disks are listed automatically. For more information about Amazon EC2 instance storage, see Instance storage in the *Amazon EC2 User Guide.*

Storage requirements

In addition to 80 GiB of disk space for the VM, you also need additional disks for your gateway.

Gateway type	Cache (minimum)	Cache (maximum)			
File Gateway	150 GiB	64 TiB			

> **Note**
>
> You can configure one or more local drives for your cache, up to the maximum capacity. When adding cache to an existing gateway, it's important to create new disks in your host (hypervisor or Amazon EC2 instance). Don't change the size of existing disks if the disks have been previously allocated as a cache.

Network and firewall requirements

Your gateway requires access to the internet, local networks, Domain Name Service (DNS) servers, firewalls, routers, and so on.

Network bandwidth requirements vary based on the quantity of data that is uploaded and downloaded by the gateway. A minimum of 100Mbps is required to successfully download, activate, and update the gateway. Your data transfer patterns will determine the bandwidth necessary to support your workload.

Following, you can find information about required ports and how to allow access through firewalls and routers.

> **Note**
>
> In some cases, you might deploy your gateway on Amazon EC2 or use other types of deployment (including on-premises) with network security policies that restrict AWS IP address ranges. In these cases, your gateway might experience service connectivity issues when the AWS IP range values changes. The AWS IP address range values that you need to use are in the Amazon service subset for the AWS Region that you activate your gateway in. For the current IP range values, see AWS IP address ranges in the *AWS General Reference*.

Topics

Port requirements

FSx File Gateway requires specific ports to be allowed through your network security for successful deployment and operation. Some ports are required for all gateways, while others are required only for specific configurations, such as when connecting to VPC endpoints.

For FSx File Gateway, you must use Microsoft Active Directory to allow domain users to access a Server Message Block (SMB) file share. You can join your File Gateway to any valid Microsoft Windows domain (resolvable by DNS).

You can also use the AWS Directory Service to create an AWS Managed Microsoft AD in the Amazon Web Services Cloud. For most AWS Managed Microsoft AD deployments, you need to configure the Dynamic Host Configuration Protocol (DHCP) service for your VPC. For information about creating

a DHCP options set, see Create a DHCP options set in the *AWS Directory Service Administration Guide*.

The following table lists the necessary ports and describes conditional requirements in the **Notes** column.

Port requirements for FSx File Gateway

Network Element	From	To	Protocol	Port	Inbound	Outboun	Required	Notes
Web browser	Your web browser	Storage Gateway VM	TCP HTTP	80	✓	✓	✓	Used by local systems to obtain the Storage Gateway activation key. Port 80 is used only during activation of a Storage Gateway appliance. A Storage Gateway VM doesn't require port 80

Network Element	From	To	Protocol	Port	Inbound	Outboun	Required	Notes
								to be publicly accessible. The required level of access to port 80 depends on your network configuration. If you activate your gateway from the Storage Gateway Management Console, the host from which you connect to the console must

Network Element	From	To	Protocol	Port	Inbound	Outboun	Required	Notes
								have access to your gateway's port 80.
Web browser	Storage Gateway VM	AWS	TCP HTTPS	443	✓	✓	✓	AWS Management Console (all other operations)
DNS	Storage Gateway VM	Domain Name Service (DNS) server	TCP & UDP DNS	53	✓	✓	✓	Used for communication between a Storage Gateway VM and the DNS server for IP name resolution.

Network Element	From	To	Protocol	Port	Inbound	Outboun	Required	Notes
NTP	Storage Gateway VM	Network Time Protocol (NTP) server	TCP & UDP NTP	123	✓	✓	✓	Used by on-premises systems to synchronize VM time to the host time. A Storage Gateway VM is configured to use the following NTP servers: • 0.amazon.pool.ntp.org • 1.amazon.pool.ntp.org • 2.amazon.pool.ntp.org

Network Element	From	To	Protocol	Port	Inbound	Outboun	Required	Notes
								• 3.amazon.pool.ntp.org **Note** Not required for gateways hosted on Amazon EC2.

Network Element	From	To	Protocol	Port	Inbound	Outboun	Required	Notes
Storage Gateway	Storage Gateway VM	Support Endpoint	TCP SSH	22	✓	✓	✓	Allows Support to access your gateway to help you with troubleshooting gateway issues. You don't need this port open for the normal operation of your gateway, but it is required for troubleshooting. For a list of support endpoints

Network Element	From	To	Protocol	Port	Inbound	Outboun	Required	Notes
								, see Support endpoints .
Storage Gateway	Storage Gateway VM	AWS	TCP HTTPS	443	✓	✓	✓	Management control
Amazon CloudFront	Storage Gateway VM	AWS	TCP HTTPS	443	✓	✓	✓	For activation
VPC	Storage Gateway VM	AWS	TCP HTTPS	443	✓	✓	✓*	Management control *Required only when using VPC endpoints
VPC	Storage Gateway VM	AWS	TCP HTTPS	1026		✓	✓*	Control Plane endpoint *Required only when using VPC endpoints

Network Element	From	To	Protocol	Port	Inbound	Outboun	Required	Notes
VPC	Storage Gateway VM	AWS	TCP HTTPS	1027		✓	✓*	Anon Control Plane (for activation) *Required only when using VPC endpoints
VPC	Storage Gateway VM	AWS	TCP HTTPS	1028		✓	✓*	Proxy endpoint *Required only when using VPC endpoints
VPC	Storage Gateway VM	AWS	TCP HTTPS	1031		✓	✓*	Data Plane *Required only when using VPC endpoints

Network Element	From	To	Protocol	Port	Inbound	Outboun	Required	Notes
VPC	Storage Gateway VM	AWS	TCP HTTPS	2222		✓	✓*	SSH Support Channel for VPCe *Required only for opening support channel when using VPC endpoints
VPC	Storage Gateway VM	AWS	TCP HTTPS	443	✓	✓	✓*	Management control *Required only when using VPC endpoints

Network Element	From	To	Protocol	Port	Inbound	Outboun	Required	Notes
File share client	SMB Client	Storage Gateway VM	TCP or UDP SMBv3	445	✓	✓	✓	File sharing data transfer session service. Replaces ports 137–139 for Microsoft Windows NT and later.
Microsoft Active Directory	Storage Gateway VM	Active Directory server	UDP NetBIOS	137	✓	✓	✓	Name service
Microsoft Active Directory	Storage Gateway VM	Active Directory server	UDP NetBIOS	138	✓	✓	✓	Datagram service
Microsoft Active Directory	Storage Gateway VM	Active Directory server	TCP & UDP LDAP	389	✓	✓	✓	Directory System Agent (DSA) client connection
Microsoft Active Directory	Storage Gateway VM	Active Directory server	TCP & UDP Kerberos	88	✓	✓	✓	Kerberos

Network Element	From	To	Protocol	Port	Inbound	Outboun	Required	Notes
Microsoft Active Directory	Storage Gateway VM	Active Directory server	TCP Distributed Computin Environm nt/End Point Mapper (DCE/ EMAP)	135	✓	✓	✓	RPC
Amazon FSx connectio n	Storage Gateway VM	FSx for Windows File Server	TCP or UDP SMBv3	445	✓	✓	✓	File sharing data transfer session service

Networking and firewall requirements for the Storage Gateway Hardware Appliance

Each Storage Gateway Hardware Appliance requires the following network services:

- **Internet access** – an always-on network connection to the internet through any network interface on the server.
- **DNS services** – DNS services for communication between the hardware appliance and DNS server.
- **Time synchronization** – an automatically configured Amazon NTP time service must be reachable.
- **IP address** – A DHCP or static IPv4 address assigned. You cannot assign an IPv6 address.

There are five physical network ports at the rear of the Dell PowerEdge R640 server. From left to right (facing the back of the server) these ports are as follows:

1. iDRAC
2. em1
3. em2
4. em3
5. em4

You can use the iDRAC port for remote server management.

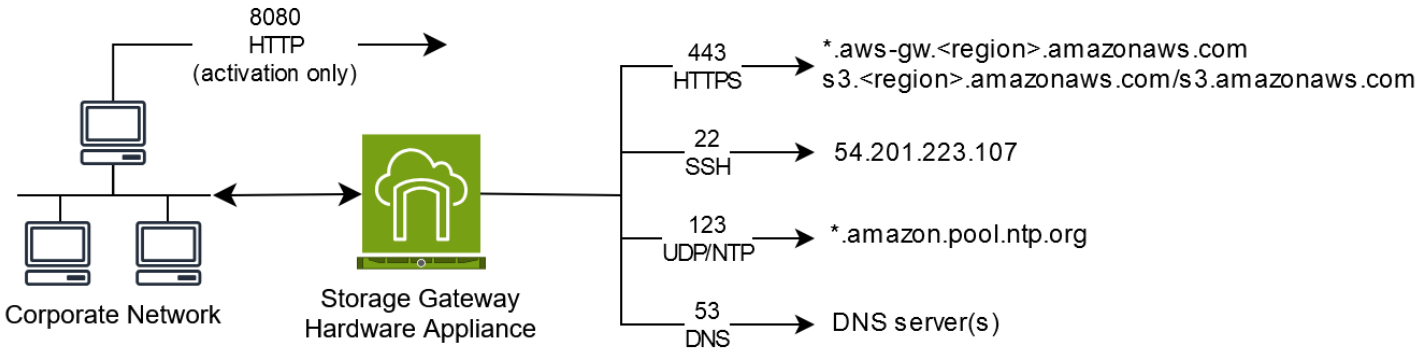

A hardware appliance requires the following ports to operate.

Protocol	Port	Direction	Source	Destination	Usage
SSH	22	Outbound	Hardware appliance	`54.201.223.107`	Support channel
DNS	53	Outbound	Hardware appliance	DNS servers	Name resolution
UDP/NTP	123	Outbound	Hardware appliance	`*.amazon.pool.ntp.org`	Time synchronization
HTTPS	443	Outbound	Hardware appliance	`*.amazonaws.com`	Data transfer

Protocol	Port	Direction	Source	Destination	Usage
HTTP	8080	Inbound	AWS	Hardware appliance	Activation (only briefly)

To perform as designed, a hardware appliance requires network and firewall settings as follows:

- Configure all connected network interfaces in the hardware console.
- Make sure that each network interface is on a unique subnet.
- Provide all connected network interfaces with outbound access to the endpoints listed in the diagram preceding.
- Configure at least one network interface to support the hardware appliance. For more information, see Configuring hardware appliance network parameters.

Note

For an illustration showing the back of the server with its ports, see Physically installing your hardware appliance.

All IP addresses on the same network interface (NIC), whether for a gateway or a host, must be on the same subnet. The following illustration shows the addressing scheme.

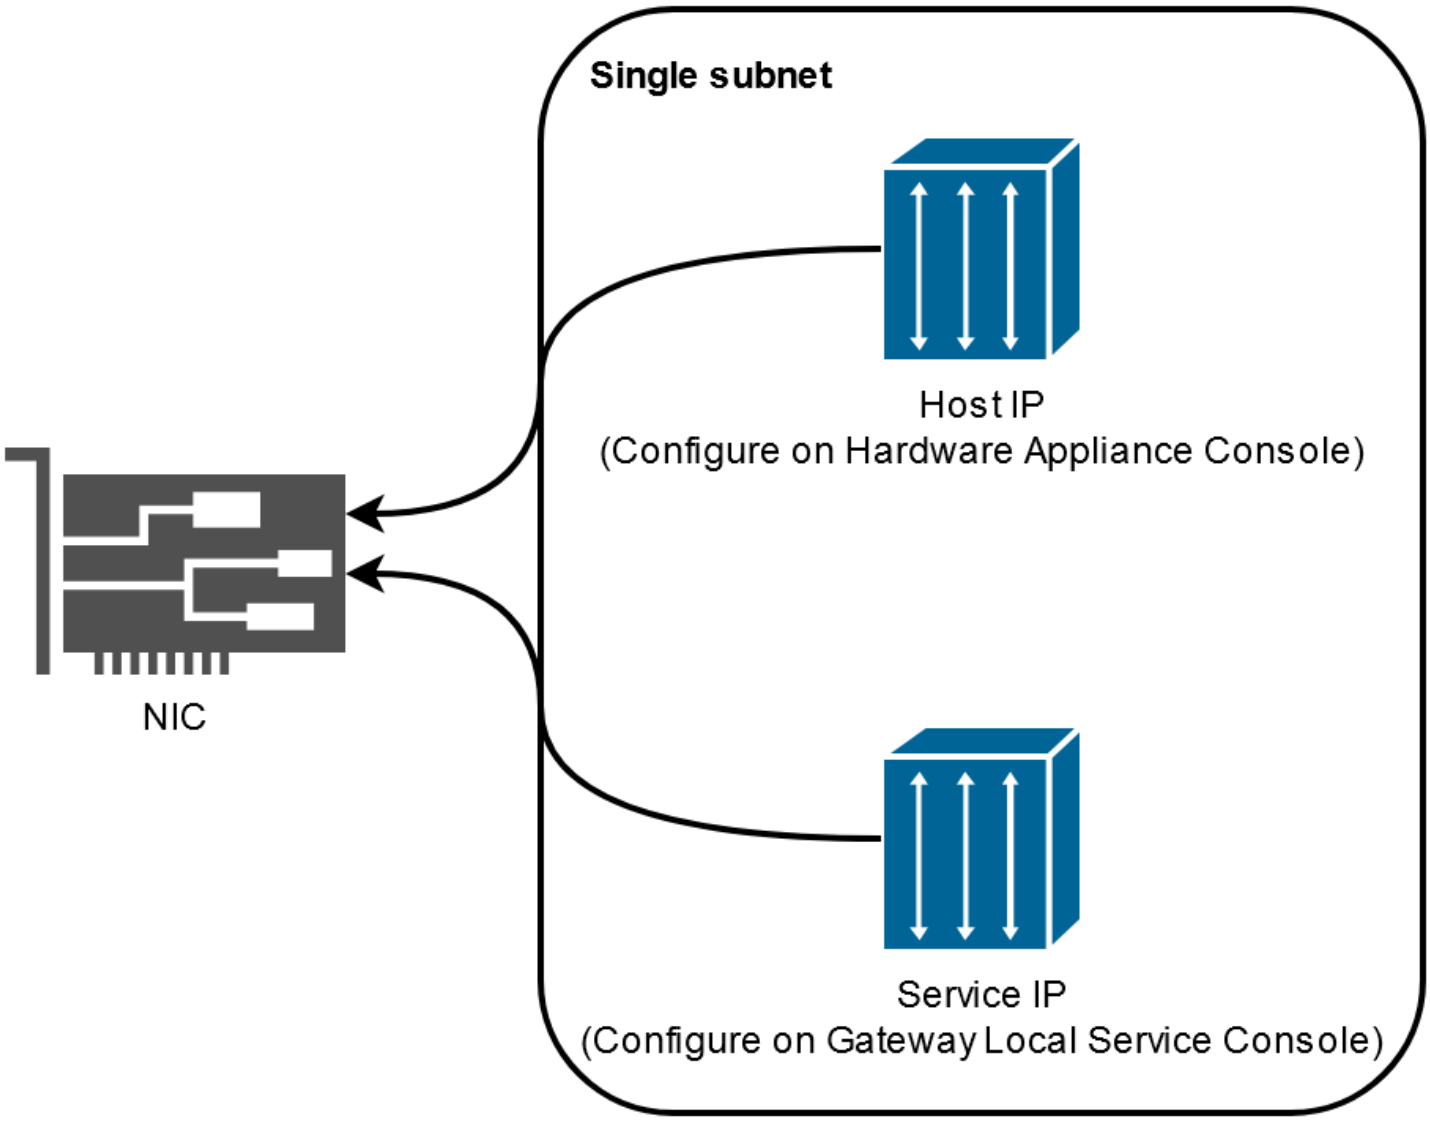

For more information about activating and configuring a hardware appliance, see Using the AWS Storage Gateway Hardware Appliance.

Allowing AWS Storage Gateway access through firewalls and routers

Your gateway requires access to the following service endpoints to communicate with AWS. If you use a firewall or router to filter or limit network traffic, you must configure your firewall and router to allow these service endpoints for outbound communication to AWS.

Note

If you configure private VPC endpoints for your Storage Gateway to use for connection and data transfer to and from AWS, your gateway does not require access to the public internet. For more information, see Activating a gateway in a virtual private cloud.

Important

Replace *region* in the following endpoint examples with the correct AWS Region string for your gateway, such as `us-west-2`.
Replace *amzn-s3-demo-bucket* with the actual name of the Amazon S3 bucket in your deployment. You can also use an asterisk (*) in place of *amzn-s3-demo-bucket* to create a wildcard entry in your firewall rules, which will allowlist the service endpoint for all bucket names.
If your gateways are deployed in AWS Regions in the United States or Canada and require Federal Information Processing Standard (FIPS) compliant endpoint connections, replace *s3* with `s3-fips`.

The following service endpoint is required by all gateways for head-bucket operations.

```
bucket-name.s3.region.amazonaws.com:443
```

The following service endpoints are required by all gateways for control path (`anon-cp`, `client-cp`, `proxy-app`) and data path (`dp-1`) operations.

```
anon-cp.storagegateway.region.amazonaws.com:443
client-cp.storagegateway.region.amazonaws.com:443
proxy-app.storagegateway.region.amazonaws.com:443
dp-1.storagegateway.region.amazonaws.com:443
```

The following gateway service endpoint is required to make API calls.

```
storagegateway.region.amazonaws.com:443
```

The following example is a gateway service endpoint in the US West (Oregon) Region (`us-west-2`).

```
storagegateway.us-west-2.amazonaws.com:443
```

In addition to the Storage Gateway and Amazon S3 service endpoints, Storage Gateway VMs also require network access to the following NTP servers:

```
0.amazon.pool.ntp.org
1.amazon.pool.ntp.org
2.amazon.pool.ntp.org
3.amazon.pool.ntp.org
```

- For a complete list of supported AWS Regions and AWS service endpoints that you can use with Storage Gateway, see AWS Storage Gateway endpoints and quotas in the *AWS General Reference*.
- For a list supported AWS Regions that you can use with the hardware appliance, see Storage Gateway hardware appliance Regions in the *AWS General Reference*.

Configuring security groups for your Amazon EC2 gateway instance

In AWS Storage Gateway, a security group controls traffic to your Amazon EC2 gateway instance. When you configure a security group, we recommend the following:

- The security group should not allow incoming connections from the outside internet. It should allow only instances within the gateway security group to communicate with the gateway.

  If you need to allow instances to connect to the gateway from outside its security group, we recommend that you allow connections only on port 80 (for activation).
- If you want to activate your gateway from an Amazon EC2 host outside the gateway security group, allow incoming connections on port 80 from the IP address of that host. If you cannot determine the activating host's IP address, you can open port 80, activate your gateway, and then close access on port 80 after completing activation.
- Allow port 22 access only if you are using Support for troubleshooting purposes. For more information, see You want Support to help troubleshoot your Amazon EC2 gateway.

Supported hypervisors and host requirements

You can run Storage Gateway on-premises as either a virtual machine (VM) appliance or a physical hardware appliance, or in AWS as an Amazon EC2 instance.

Storage Gateway supports the following hypervisor versions and hosts:

- VMware ESXi Hypervisor (version 7.0 or 8.0) – For this setup, you also need a VMware vSphere client to connect to the host.
- Microsoft Hyper-V Hypervisor (version 2012 R2, 2016, 2019, or 2022) – A free, standalone version of Hyper-V is available at the Microsoft Download Center. For this setup, you need a Microsoft Hyper-V Manager on a Microsoft Windows client computer to connect to the host.
- Linux Kernel-based Virtual Machine (KVM) – A free, open-source virtualization technology. KVM is included in all versions of Linux version 2.6.20 and newer. Storage Gateway is tested and supported for the CentOS/RHEL 7.7, RHEL 8.6 Ubuntu 16.04 LTS, and Ubuntu 18.04 LTS distributions. Any other modern Linux distribution may work, but function or performance is not guaranteed. We recommend this option if you already have a KVM environment up and running and you are already familiar with how KVM works.
- Amazon EC2 instance – Storage Gateway provides an Amazon Machine Image (AMI) that contains the gateway VM image. For information about how to deploy a gateway on Amazon EC2, see Deploy a default Amazon EC2 host for FSx File Gateway.
- Storage Gateway Hardware Appliance – Storage Gateway provides a physical hardware appliance as an on-premises deployment option for locations with limited virtual machine infrastructure.

> **Note**
>
> Storage Gateway doesn't support recovering a gateway from a VM that was created from a snapshot or clone of another gateway VM or from your Amazon EC2 AMI. If your gateway VM malfunctions, activate a new gateway and recover your data to that gateway. For more information, see Recovering from an unexpected virtual machine shutdown.
> Storage Gateway doesn't support dynamic memory and virtual memory ballooning.

Supported SMB clients for File Gateway

File Gateway supports the following Service Message Block (SMB) clients:

- Microsoft Windows Server 2008 R2 and later
- Windows desktop versions: 10, 8, and 7.
- Windows Terminal Server running on Windows Server 2008 and later

> **Note**
>
> Server Message Block encryption requires clients that support SMB v3.x dialects.

Supported file system operations for File Gateway

Your SMB client can write, read, delete, and truncate files. When clients send writes to Storage Gateway, it writes to local cache synchronously. Then it writes to Amazon FSx asynchronously through optimized transfers. Reads are first served through the local cache. If data is not available, it's fetched through Amazon FSx as a read-through cache.

Writes and reads are optimized in that only the parts that are changed or requested are transferred through your gateway. Deletes remove files from Amazon FSx.

Managing local disks for your gateway

The gateway virtual machine (VM) uses the local disks that you allocate on-premises for buffering and storage. A File Gateway that you create on an Amazon EC2 instance will use Amazon EBS volumes as local disks. The number and size of disks that you want to allocate for your gateway is up to you. The gateway uses the cache storage that you allocate to provide low-latency access to your recently accessed data. The cache storage acts as the on-premises durable store for data that is pending upload to Amazon FSx. File Gateways require at least one 150 GiB disk to use as a cache. After the initial configuration and deployment of your gateway, you can add more disks for cache storage as your workload demands increase. This section contains the following topics, which describe concepts and procedures related to managing local disks.

Topics

- Deciding the amount of local disk storage - Learn how to determine the number and size of local cache disks to allocate for your File Gateway.
- Configuring additional cache storage - Learn how to increase the cache storage capacity of your File Gateway as your application needs change.
- Using ephemeral storage with EC2 gateways - Learn how to prevent data loss when using ephemeral disk storage with File Gateway.

Deciding the amount of local disk storage

When deploying an FSx File Gateway, consider how much cache disk to allocate. FSx File Gateway uses a least recently used algorithm to automatically evict data from the cache. The cache on an FSx File Gateway is shared between all of the file shares on that gateway. If you have multiple active shares, it's important to note that heavy utilization on one share could impact the amount of cache resources that another share has access to, possibly impacting performance.

When determining how much cache disk you need for a given workload, it's important to note that you can always add cache disk to your gateway (up to the current quotas on FSx File Gateway), but you can't decrease the cache for a given gateway. You can perform a basic analysis on the dataset to determine the right amount of cache disk, but there's not a way to determine exactly how much data is 'hot,' and needs to be stored locally, versus 'cold' and can be tiered to the cloud. Workloads change over time, and FSx File Gateway provides flexibility and elasticity related to the amount of resources that can be consumed. The amount of cache can always be increased, so starting small and increasing as needed is often the most cost-effective approach.

You can use an initial approximation of 150 GiB to provision disks for the cache storage during gateway setup. You can then use Amazon CloudWatch operational metrics to monitor the cache storage usage and provision more storage as needed using the console. For information on using the metrics and setting up alarms, see Performance and optimization.

Note

Underlying physical storage resources are represented as a data store in VMware. When you deploy the gateway VM, you choose a data store on which to store the VM files. When you provision a local disk (for example, to use as cache storage), you have the option to store the virtual disk in the same data store as the VM or a different data store.
If you have more than one data store, we strongly recommend that you choose one data store for the cache storage. A data store that is backed by only one underlying physical disk can lead to poor performance in some situations when it is used to back both the cache storage. This is also true if the backup is a less-performant RAID configuration such as RAID1.

Configuring additional cache storage

As your application needs change, you can increase the gateway's cache storage capacity. You can add storage capacity to your gateway without interrupting functionality or causing downtime. When you add more storage, you do so with the gateway VM turned on.

> **Important**
>
> When adding cache to an existing gateway, you must create new disks on the gateway host hypervisor or Amazon EC2 instance. Do not remove or change the size of existing disks that have already been allocated as cache.

To configure additional cache storage for your gateway

1. Provision one or more new disks on your gateway host hypervisor or Amazon EC2 instance. For information about how to provision a disk on a hypervisor, see your hypervisor's documentation. For information about provisioning Amazon EBS volumes for an Amazon EC2 instance, see Amazon EBS volumes in the *Amazon Elastic Compute Cloud User Guide for Linux Instances*. In the following steps, you will configure this disk as cache storage.
2. Open the Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
3. In the navigation pane, choose **Gateways**.
4. Search for your gateway and select it from the list.
5. From the **Actions** menu, choose **Configure cache storage**.
6. In the **Configure cache storage** section, identify the disks you provisioned. If you don't see your disks, choose the refresh icon to refresh the list. For each disk, choose **Cache** from the **Allocated to** drop-down menu.

   > **Note**
   >
   > **Cache** is the only available option for allocating disks on a File Gateway.

7. Choose **Save changes** to save your configuration settings.

Using ephemeral storage with EC2 gateways

We do not recommend the use of ephemeral disks for cache storage on FSx File Gateways.

Ephemeral disks provide temporary block-level storage for your Amazon EC2 instance. When you launch your gateway with an Amazon EC2 Amazon Machine Image and the instance type you select supports ephemeral storage, the ephemeral disks are listed automatically. You can select one of the disks to store your gateway's cache data. For more information, see Amazon EC2 instance store in the *Amazon EC2 User Guide*.

Data that applications write to the gateway is stored synchronously in cache on the ephemeral disks, and then asynchronously uploaded to durable storage in FSx for Windows File Server. If the Amazon EC2 instance is stopped after data is written to ephemeral storage, but before an asynchronous upload occurs, any data that has not yet been uploaded to FSx for Windows File Server can be lost.

Important

If you stop and start an Amazon EC2 gateway that uses ephemeral storage, the gateway will be permanently offline. This happens because the physical storage disk is replaced. There is no work-around for this issue. The only resolution is to delete the gateway and activate a new one on a new EC2 instance.

Using the AWS Storage Gateway Hardware Appliance

> **Note**
>
> End of availability notice: As of May 12, 2025, the AWS Storage Gateway Hardware Appliance will no longer be offered. Existing customers with the AWS Storage Gateway Hardware Appliance can continue to use and receive support until May 2028. As an alternative, you can use the AWS Storage Gateway service to give your applications on-premises and in-cloud access to virtually unlimited cloud storage.

The AWS Storage Gateway Hardware Appliance is a physical hardware appliance with the Storage Gateway software preinstalled on a validated server configuration. You can manage the hardware appliances in your deployment from the **Hardware appliance overview** page in the AWS Storage Gateway console.

The hardware appliance is a high-performance 1U server that you can deploy in your data center, or on-premises inside your corporate firewall. When you buy and activate your hardware appliance, the activation process associates the hardware appliance with your AWS account. After activation, your hardware appliance appears in the console on the **Hardware appliance overview** page. You can configure the hardware appliance as an S3 File Gateway, FSx File Gateway, Tape Gateway, or Volume Gateway type. The procedure that you use to deploy these gateway types on a hardware appliance is same as on a virtual platform.

For a list of supported AWS Regions where the AWS Storage Gateway Hardware Appliance is available for activation and use, see AWS Storage Gateway Hardware Appliance Regions in the *AWS General Reference*.

In the sections that follow, you can find instructions about how to set up, rack mount, power, configure, activate, launch, use, and delete an AWS Storage Gateway Hardware Appliance.

Topics

- Creating a gateway on your hardware appliance
- Configuring a gateway IP address on the hardware appliance
- Removing gateway software from your hardware appliance
- Deleting your AWS Storage Gateway Hardware Appliance

Setting up your AWS Storage Gateway Hardware Appliance

> **Note**
>
> End of availability notice: As of May 12, 2025, the AWS Storage Gateway Hardware Appliance will no longer be offered. Existing customers with the AWS Storage Gateway Hardware Appliance can continue to use and receive support until May 2028. As an alternative, you can use the AWS Storage Gateway service to give your applications on-premises and in-cloud access to virtually unlimited cloud storage.

After you receive your Storage Gateway Hardware Appliance, you use the hardware appliance local console to configure networking to provide an always-on connection to AWS and activate your appliance. Activation associates your appliance with the AWS account that is used during the activation process. After the appliance is activated, you can launch an S3 File Gateway, FSx File Gateway, Tape Gateway, or Volume Gateway from the Storage Gateway console.

To install and configure your hardware appliance

1. Rack-mount the appliance, and plug in power and network connections. For more information, see Physically installing your hardware appliance.
2. Set the Internet Protocol version 4 (IPv4) addresses for the hardware appliance (the host). For more information, see Configuring hardware appliance network parameters.
3. Activate the hardware appliance on the console **Hardware appliance overview** page in the AWS Region of your choice. For more information, see Activating your AWS Storage Gateway Hardware Appliance.
4. Create a gateway on your hardware appliance. For more information, see Creating your gateway.

   You set up gateways on your hardware appliance the same way that you set up gateways on VMware ESXi, Microsoft Hyper-V, Linux Kernel-based Virtual Machine (KVM), or Amazon EC2.

Increasing the usable cache storage

You can increase the usable storage on the hardware appliance from 5 TB to 12 TB. Doing this provides a larger cache for low latency access to data in AWS. If you ordered the 5 TB model, you can increase the usable storage to 12 TB by buying five 1.92 TB SSDs (solid state drives).

You can then add them to the hardware appliance before you activate it. If you have already activated the hardware appliance and want to increase the usable storage on the appliance to 12 TB, do the following:

1. Reset the hardware appliance to its factory settings. Contact AWS Support for instructions on how to do this.
2. Add five 1.92 TB SSDs to the appliance.

Network interface card options

Depending on the model of appliance you ordered, it may come with a 10G-Base-T RJ45 copper, or a 10G DA/SFP+ network card.

- 10G-Base-T NIC configuration:
  - Use CAT6 cables for 10G or CAT5(e) for 1G
- 10G DA/SFP+ NIC configuration:
  - Use Twinax copper Direct Attach Cables up to 5 meters
  - Dell/Intel compatible SFP+ optical modules (SR or LR)
  - SFP/SFP+ copper transceiver for 1G-Base-T or 10G-Base-T

Physically installing your hardware appliance

> **Note**
>
> End of availability notice: As of May 12, 2025, the AWS Storage Gateway Hardware Appliance will no longer be offered. Existing customers with the AWS Storage Gateway Hardware Appliance can continue to use and receive support until May 2028. As an alternative, you can use the AWS Storage Gateway service to give your applications on-premises and in-cloud access to virtually unlimited cloud storage.

Your appliance has a 1U form factor and fits in a standard International Electrotechnical Commission (IEC) compliant 19-inch rack.

Prerequisites

To install your hardware appliance, you need the following components:

- Power cables: one required, two recommended.
- Supported network cabling (depending on which Network Interface Card (NIC) is included in the hardware appliance). Twinax Copper DAC, SFP+ optical module (Intel compatible) or SFP to Base-T copper transceiver.
- Keyboard and monitor, or a keyboard, video, and mouse (KVM) switch solution.

> **Note**
>
> Before you perform the following procedure, make sure that you meet all of the requirements for the Storage Gateway Hardware Appliance as described in Networking and firewall requirements for the Storage Gateway Hardware Appliance.

To physically install your hardware appliance

1. Unbox your hardware appliance and follow the instructions contained in the box to rack-mount the server.

   The following image shows the back of the hardware appliance with ports for connecting power, ethernet, monitor, USB keyboard, and iDRAC.
   hardware appliance one rear with network and power connector labels.

   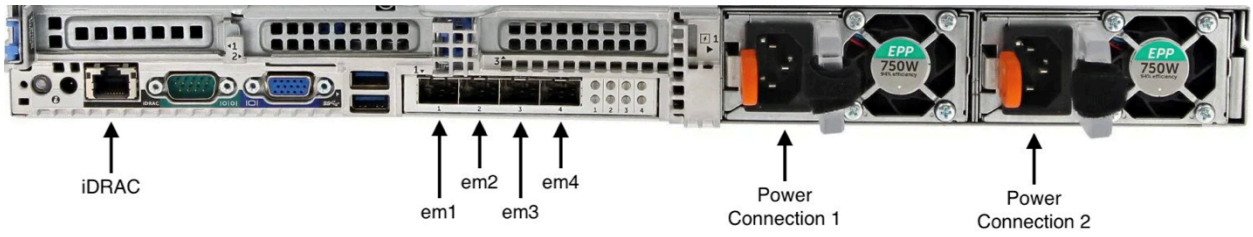

   hardware appliance one rear with network and power connector labels.
2. Plug in a power connection to each of the two power supplies. It's possible to plug in to only one power connection, but we recommend power connections to both power supplies for redundancy.

3. Plug an Ethernet cable into the em1 port to provide an always-on internet connection. The em1 port is the first of the four physical network ports on the rear, from left to right.

   > **Note**
   >
   > The hardware appliance doesn't support VLAN trunking. Set up the switch port to which you are connecting the hardware appliance as a non-trunked VLAN port.

4. Plug in the keyboard and monitor.

5. Power on the server by pressing the **Power** button on the front panel, as shown in the following image.
   hardware appliance front with power button label.

   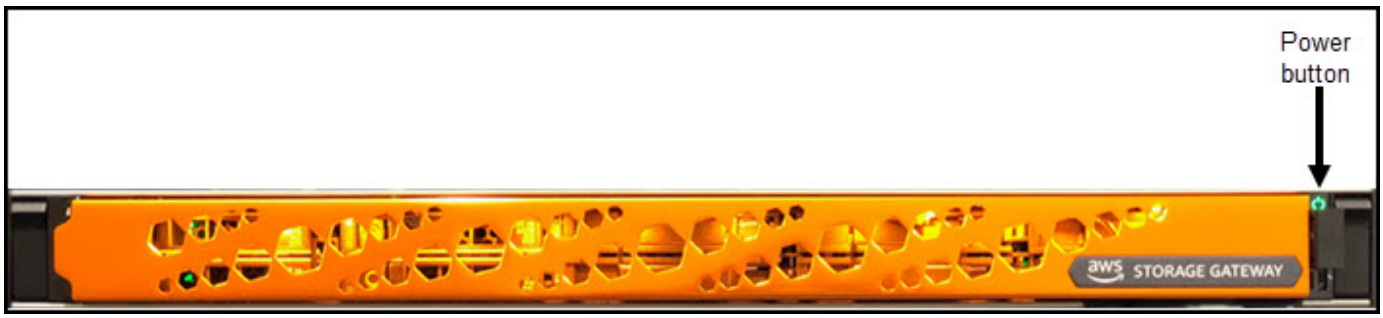

   hardware appliance front with power button label.

Next step

Accessing the hardware appliance console

Accessing the hardware appliance console

> **Note**
>
> End of availability notice: As of May 12, 2025, the AWS Storage Gateway Hardware Appliance will no longer be offered. Existing customers with the AWS Storage Gateway Hardware Appliance can continue to use and receive support until May 2028. As an alternative, you can use the AWS Storage Gateway service to give your applications on-premises and in-cloud access to virtually unlimited cloud storage.

When you power on your hardware appliance, the hardware appliance console appears on the monitor. The hardware appliance console presents a user interface specific to AWS that you can use to set an administrator password, configure initial network parameters, and open a support channel to AWS.

To work with the hardware appliance console, enter text from the keyboard and use the Up, Down, Right, and Left Arrow keys to move about the screen in the indicated direction. Use the Tab key to move forward in order through items on-screen. On some setups, you can use the Shift +Tab keystroke to move sequentially backward. Use the Enter key to save selections, or to choose a button on the screen.

The first time the hardware appliance console appears, the **Welcome** page is displayed, and you are prompted to set a password for the *admin* user account before you can access the console.

To set an admin password

- At the **Please set your login password** prompt, do the following:

  a. For **Set Password**, enter a password, and then press Down arrow.

  b. For **Confirm**, re-enter your password, and then choose **Save Password**.

After you set your password, the hardware console **Home** page appears. The **Home** page displays network information for the **em1**, **em2**, **em3**, and **em4** network interfaces, and has the following menu options:

- Configure Network
- Open Service Console
- Change Password
- Logout
- Open Support Console

Next step

Configuring hardware appliance network parameters

Configuring hardware appliance network parameters

> **Note**
>
> End of availability notice: As of May 12, 2025, the AWS Storage Gateway Hardware Appliance will no longer be offered. Existing customers with the AWS Storage Gateway Hardware Appliance can continue to use and receive support until May 2028. As an alternative, you can use the AWS Storage Gateway service to give your applications on-premises and in-cloud access to virtually unlimited cloud storage.

After the hardware appliance boots up and you set your admin user password in the hardware console as described in Accessing the hardware appliance console, use the following procedure to configure network parameters so your hardware appliance can connect to AWS.

To set a network address

1. From the **Home** page, choose **Configure Network** and then press `Enter`. The **Configure Network** page appears. The **Configure Network** page shows IP and DNS information for each of the 4 network interfaces on the hardware appliance, and includes menu options to configure **DHCP** or **Static** addresses for each.
2. For the **em1** interface, do one of the following:

   - Choose **DHCP** and press `Enter` to use the IPv4 address assigned by your Dynamic Host Configuration Protocol (DHCP) server to your physical network port.

     Note this address for later use in the activation step.
   - Choose **Static** and press `Enter` to configure a static IPv4 address.

     Enter a valid **IP Address**, **Subnet Mask**, **Gateway**, and **DNS** server address for the **em1** network interface.

     When finished, choose **Save** and then press `Enter` to save the configuration.

> **Note**
>
> You can use this procedure to configure other network interfaces in addition to **em1**.
> If you configure other interfaces, they must provide the same always-on connection to the AWS endpoints listed in the requirements.
> Network bonding and Link Aggregation Control Protocol (LACP) are not supported by the hardware appliance or by Storage Gateway.
> We do not recommend configuring multiple network interfaces on the same subnet as this can sometimes cause routing issues.

To log out of the hardware console

1. Choose **Back** and press `Enter` to return to the **Home** page.
2. Choose **Logout** and press `Enter` to return to the **Welcome** page.

Next step

Activating your AWS Storage Gateway Hardware Appliance

Activating your AWS Storage Gateway Hardware Appliance

> **Note**
>
> End of availability notice: As of May 12, 2025, the AWS Storage Gateway Hardware Appliance will no longer be offered. Existing customers with the AWS Storage Gateway Hardware Appliance can continue to use and receive support until May 2028. As an alternative, you can use the AWS Storage Gateway service to give your applications on-premises and in-cloud access to virtually unlimited cloud storage.

After configuring your IP address, you enter this IP address on the **Hardware** page of the AWS Storage Gateway console to activate your hardware appliance. The activation process registers the appliance to your AWS account.

You can choose to activate your hardware appliance in any of the supported AWS Regions. For a list of supported AWS Regions, see Storage Gateway Hardware Appliance Regions in the *AWS General Reference*.

To activate your AWS Storage Gateway Hardware Appliance

1. Open the AWS Storage Gateway Management Console and sign in with the account credentials you want to use to activate your hardware.

   > **Note**
   >
   > For activation only, the following must be true:
   >
   > - Your browser must be on the same network as your hardware appliance.
   > - Your firewall must allow HTTP access on port 8080 to the appliance for inbound traffic.

2. Choose **Hardware** from the navigation menu on the left side of the page.
3. Choose **Activate appliance**.
4. For **IP Address**, enter the IP address that you configured for your hardware appliance, then choose **Connect**.

   For more information about configuring the IP address, see Configuring network parameters.
5. For **Name**, enter a name for your hardware appliance. Names can be up to 255 characters long and can't include a slash character.
6. For **Hardware appliance time zone**, enter the local time zone from which most of the workload for the gateway will be generated., then choose **Next**.

   The time zone controls when hardware updates take place, with 2 a.m. used as the default scheduled time to perform updates. Ideally, if the time zone is set properly, updates will take place outside of the local working day window by default.
7. Review the activation parameters in the Hardware appliance detail section. You can choose **Previous** to go back and make changes if necessary. Otherwise, choose **Activate** to finish the activation.

A banner appears on the **Hardware appliance overview** page, indicating that the hardware appliance has been successfully activated.

At this point, the appliance is associated with your account. The next step is to configure and launch an S3 File Gateway, FSx File Gateway, Tape Gateway, or Volume Gateway on the new appliance.

Next step

Creating a gateway on your hardware appliance

Creating a gateway on your hardware appliance

> **Note**
>
> End of availability notice: As of May 12, 2025, the AWS Storage Gateway Hardware Appliance will no longer be offered. Existing customers with the AWS Storage Gateway Hardware Appliance can continue to use and receive support until May 2028. As an alternative, you can use the AWS Storage Gateway service to give your applications on-premises and in-cloud access to virtually unlimited cloud storage.

You can create an S3 File Gateway, FSx File Gateway, Tape Gateway, or Volume Gateway on any AWS Storage Gateway Hardware Appliance in your deployment.

To create a gateway on your hardware appliance

1. Sign in to the AWS Management Console and open the Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. Follow the procedures described in Creating Your Gateway to set up, connect, and configure the type of Storage Gateway that you want to deploy.

When you finish creating your gateway in the Storage Gateway console, the Storage Gateway software automatically starts installing on the hardware appliance. If you use Dynamic Host Configuration Protocol (DHCP), it can take 5 to 10 minutes for a gateway to display as online in the console. To assign a static IP address to your installed gateway, see Configuring an IP address for the gateway.

To assign a static IP address to your installed gateway, you next configure the gateway's network interfaces so your applications can use it.

Next step

Configuring a gateway IP address on the hardware appliance

Configuring a gateway IP address on the hardware appliance

> **Note**
>
> End of availability notice: As of May 12, 2025, the AWS Storage Gateway Hardware Appliance will no longer be offered. Existing customers with the AWS Storage Gateway Hardware Appliance can continue to use and receive support until May 2028. As an alternative, you can use the AWS Storage Gateway service to give your applications on-premises and in-cloud access to virtually unlimited cloud storage.

Before you activated your hardware appliance, you assigned an IP address to its physical network interface. Now that you have activated the appliance and launched your Storage Gateway on it, you need to assign another IP address to the Storage Gateway virtual machine that runs on the hardware appliance. To assign a static IP address to a gateway installed on your hardware appliance, configure the IP address from the gateway local console for that gateway. Your applications (such as your NFS or SMB client) connect to this IP address. You can access the gateway local console from the hardware appliance console using the **Open Service Console** option.

To configure an IP address on your appliance to work with applications

1. On the hardware console, choose **Open Service Console** and then press `Enter` to open the login page for the gateway local console.
2. The AWS Storage Gateway local console login page prompts you to login to change your network configuration and other settings.

   The default account is `admin` and the default password is `password`.

   > **Note**
   >
   > We recommend changing the default password by entering the corresponding numeral for **Gateway Console** from the **AWS Appliance Activation - Configuration** main menu, then running the `passwd` command. For information about how to run the command,

see Running Storage Gateway commands on the local console. You can also set the password from the Storage Gateway console. For more information, see Setting the local console password from the Storage Gateway console .

3. The **AWS Appliance Activation - Configuration** page includes the following menu options:

   - HTTP/SOCKS Proxy Configuration
   - Network Configuration
   - Test Network Connectivity
   - View System Resource Check
   - System Time Management
   - License Information
   - Command Prompt

   > **Note**
   >
   > Some options appear only for specific gateway types or host platforms.

   Enter the corresponding numeral to navigate to the **Network Configuration** page.

4. Do one of the following to configure the gateway IP address:

   - To use the IP address assigned by your Dynamic Host Configuration Protocol (DHCP) server, enter the corresponding numeral for **Configure DHCP**, and then enter valid DHCP configuration information on the following page.
   - To assign a static IP address, enter the corresponding numeral for **Configure Static IP**, and then enter valid IP address and DNS information on the following page.

     > **Note**
     >
     > The IP address you specify here must be on the same subnet as the IP address used during hardware appliance activation.

To exit the gateway local console

- Press the `Crtl+]` (close bracket) keystroke. The hardware console appears.

> **Note**
>
> The keystroke preceding is the only way to exit the gateway local console.

After your hardware appliance has been activated and configured, your appliance appears in the console. Now you can continue the setup and configuration procedure for your gateway in the Storage Gateway console. For instructions, see Configure your Amazon FSx File Gateway.

Removing gateway software from your hardware appliance

> **Note**
>
> End of availability notice: As of May 12, 2025, the AWS Storage Gateway Hardware Appliance will no longer be offered. Existing customers with the AWS Storage Gateway Hardware Appliance can continue to use and receive support until May 2028. As an alternative, you can use the AWS Storage Gateway service to give your applications on-premises and in-cloud access to virtually unlimited cloud storage.

If you no longer need a specific Storage Gateway that you have deployed on a hardware appliance, you can remove the gateway software from the hardware appliance. After you remove the gateway software, you can choose to deploy a new gateway in its place, or delete the hardware appliance itself from the Storage Gateway console. To remove gateway software from your hardware appliance, use the following procedure.

To remove a gateway from a hardware appliance

1. Open the Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. Choose **Hardware** from the navigation pane on the left side of the console page, and then choose the **Hardware appliance name** for the appliance from which you want to remove gateway software.
3. From the **Actions** drop down menu, choose **Remove gateway**.

The confirmation dialog box appears.

4. Verify that you want to remove the gateway software from the specified hardware appliance, and then type the word `remove` in the confirmation box.
5. Choose **Remove** to permanently remove the gateway software.

> **Note**
>
> After you remove the gateway software, you can't undo the action. For certain gateway types, you can lose data on deletion, particularly cached data. For more information on deleting a gateway, see Deleting your gateway and removing associated resources.

Removing the gateway doesn't delete the hardware appliance from the console. The hardware appliance remains for future gateway deployments.

Deleting your AWS Storage Gateway Hardware Appliance

> **Note**
>
> End of availability notice: As of May 12, 2025, the AWS Storage Gateway Hardware Appliance will no longer be offered. Existing customers with the AWS Storage Gateway Hardware Appliance can continue to use and receive support until May 2028. As an alternative, you can use the AWS Storage Gateway service to give your applications on-premises and in-cloud access to virtually unlimited cloud storage.

If you no longer need an AWS Storage Gateway Hardware Appliance that you have already activated, you can delete the appliance completely from your AWS account.

> **Note**
>
> To move your appliance to a different AWS account or AWS Region, you must first delete it using the following procedure, then open the gateway's support channel and contact Support to perform a soft reset. For more information, see Turning on Support access to help troubleshoot your gateway hosted on-premises.

To delete your hardware appliance

1. If you have installed a gateway on the hardware appliance, you must first remove the gateway before you can delete the appliance. For instructions on how to remove a gateway from your hardware appliance, see Removing gateway software from your hardware appliance.
2. On the Hardware page of the Storage Gateway console, choose the hardware appliance you want to delete.
3. For **Actions**, choose **Delete Appliance**. The confirmation dialog box appears.
4. Verify that you want to delete the specified hardware appliance, then type the word *delete* in the confirmation box and choose **Delete**.

   When you delete the hardware appliance, all resources associated with the gateway that is installed on the appliance are deleted, but the data on the hardware appliance itself is not deleted.

Creating your gateway

The overview sections on this page provide a high-level synopsis of how the Storage Gateway creation process works. For step-by-step procedures to create a specific type of gateway using the Storage Gateway console, see the following topics:

- Create and activate an Amazon S3 File Gateway
- Create and activate an Amazon FSx File Gateway
- Create and activate a Tape Gateway
- Create and activate a Volume Gateway

Important

Amazon FSx File Gateway is no longer available to new customers. Existing customers of FSx File Gateway can continue to use the service normally. For capabilities similar to FSx File Gateway, visit this blog post.

Overview - Gateway Activation

Gateway activation involves setting up your gateway, connecting it to AWS, then reviewing your settings and activating it.

Set up gateway

To set up your Storage Gateway, you first choose the type of gateway you want to create and the host platform on which you will run the gateway virtual appliance. You then download the gateway virtual appliance template for the platform of your choice and deploy it in your on-premises environment. You can also deploy your Storage Gateway as a physical hardware appliance that you order from your preferred reseller, or as an Amazon EC2 instance in your AWS cloud environment. When you deploy the gateway appliance, you allocate local physical disk space on the virtualization host.

Connect to AWS

The next step is to connect your gateway to AWS. To do this, you first choose the type of service endpoint you want to use for communications between the gateway virtual appliance and AWS

services in the cloud. This endpoint can be accessible from the public internet, or only from within your Amazon VPC, where you have full control over the network security configuration. You then specify the gateway's IP address or its activation key, which you can obtain by connecting to the local console on the gateway appliance.

Review and activate

At this point, you'll have an opportunity to review the gateway and connection options you chose, and make changes if necessary. When everything is set up the way you want you can activate the gateway. Before you can start using your activated gateway, you will need to configure some additional settings and create your storage resources.

Overview - Gateway Configuration

After you activate your Storage Gateway, you need to perform some additional configuration. In this step, you allocate the physical storage you provisioned on the gateway host platform to be used as either the cache or the upload buffer by the gateway appliance. You then configure settings to help monitor the health of your gateway using Amazon CloudWatch Logs and CloudWatch alarms, and add tags to help identify the gateway, if desired. Before you can start using your activated and configured gateway, you will need to create your storage resources.

Overview - Storage Resources

After you activate and configure your Storage Gateway, you need to create cloud storage resources for it to use. Depending on the type of gateway you created, you will use the Storage Gateway console to create Volumes, Tapes, or Amazon S3 or Amazon FSx files shares to associate with it. Each gateway type uses its respective resources to emulate the related type of network storage infrastructure, and transfers the data you write to it into the AWS cloud.

Create an Amazon FSx for Windows File Server file system

To create an Amazon FSx File Gateway in AWS Storage Gateway, the first step is to create an Amazon FSx for Windows File Server file system. If you've already created an Amazon FSx file system, go to the next step, Create and activate an Amazon FSx File Gateway.

Note

The following limitations apply when writing to an Amazon FSx file system from an FSx File Gateway:

- Your Amazon FSx file system and your FSx File Gateway must be owned by the same AWS account and located in the same AWS Region.
- Each gateway can support five attached file systems. When attaching a file system, the Storage Gateway console notifies you if the selected gateway is at capacity. In that case, you must choose a different gateway or detach a file system before you can attach another one.
- FSx File Gateway supports soft storage quotas (issuing warnings when users surpass their data limits), but does not support hard quotas (enforcing data limits by denying write access). Soft quotas are supported for all users except the Amazon FSx admin user. For more information about setting up storage quotas, see Storage quotas in the *Amazon FSx for Windows File Server User Guide*.
- We don't recommend using Microsoft Distributed File System (DFS) to redirect users to your Amazon FSx file system through FSx File Gateway. Instead, configure DFS to redirect directly to the Amazon FSx file system in the AWS Cloud as described in Grouping multiple file systems with DFS Namespaces in the *Amazon FSx for Windows File Server User Guide*.
- Some file operations on the FSx File Gateway, such as top-level folder renames or permission changes, can result in multiple file operations that lead to a high I/O load on your FSx for Windows File Server file system. If your file system doesn't have enough performance resources for your workload, the file system might delete shadow copies because it prioritizes availability for ongoing I/O over historical shadow copy retention.

  In the Amazon FSx console, check the **Monitoring and performance** page to see if your file system is under-provisioned. If it is, you can switch to SSD storage, increase throughput capacity, or increase SSD IOPS to handle your workload.

To create an FSx for Windows File Server file system

1. Open the AWS Management Console at https://console.aws.amazon.com/fsx/home/, and choose the Region that you want to create your gateway in.

2. Follow the instructions in Getting Started with Amazon FSx in the *Amazon FSx for Windows File Server User Guide*.

Create and activate an Amazon FSx File Gateway

In this section, you can find instructions on how to create, deploy, and activate a File Gateway in AWS Storage Gateway.

Topics

Set up an Amazon FSx File Gateway

To set up a new FSx File Gateway

1. Open the AWS Management Console at https://console.aws.amazon.com/storagegateway/home/, and choose the AWS Region where you want to create your gateway.
2. Choose **Create gateway** to open the **Set up gateway** page.
3. In the **Gateway settings** section, do the following:

    a. For **Gateway name**, enter a name for your gateway. After your gateway is created, you can search for this name to find your gateway on the list pages in the AWS Storage Gateway console.

    b. For **Gateway time zone**, choose the local time zone for the part of the world where you want to deploy your gateway.
4. In the **Gateway options** section, for **Gateway type**, choose **Amazon FSx File Gateway**.
5. In the **Platform options** section, do the following:

    a. For **Host platform**, choose the platform on which you want to deploy your gateway. Then follow the platform-specific instructions displayed on the Storage Gateway console page to set up your host platform. You can choose from the following options:

- **VMware ESXi** – Download, deploy, and configure the gateway virtual machine using VMware ESXi.
- **Microsoft Hyper-V** – Download, deploy, and configure the gateway virtual machine using Microsoft Hyper-V.
- **Linux KVM** – Download, deploy, and configure the gateway virtual machine using Linux Kernel-based Virtual Machine (KVM).
- **Amazon EC2** – Configure and launch an Amazon EC2 instance to host your gateway.
- **Hardware appliance** – Order a dedicated physical hardware appliance from AWS to host your gateway.

b. For **Confirm set up gateway**, select the check box to confirm that you performed the deployment steps for the host platform you chose. This step is not applicable for the **Hardware appliance** host platform.

6. Now that your gateway is set up, you must choose how you want it to connect and communicate with AWS. Choose **Next** to proceed.

Connect your Amazon FSx File Gateway to AWS

To connect a new FSx File Gateway to AWS

1. If you have not done so already, complete the procedure described in Set up an Amazon FSx File Gateway. When finished, choose **Next** to open the **Connect to AWS** page in the AWS Storage Gateway console.
2. In the **Endpoint options** section, for **Service endpoint**, choose the type of endpoint your gateway will use to communicate with AWS. You can choose from the following options:

   - **Publicly accessible** – Your gateway communicates with AWS over the public internet. If you select this option, use the **FIPS enabled endpoint** check box to specify whether the connection must comply with Federal Information Processing Standards (FIPS).

   > **Note**
   >
   > If you require FIPS 140-2 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS-compliant endpoint. For more information, see Federal Information Processing Standard (FIPS) 140-2.

The FIPS service endpoint is available only in some AWS Regions. For more information, see AWS Storage Gateway endpoints and quotas in the *AWS General Reference*.

- **VPC hosted** – Your gateway communicates with AWS through a private connection with your virtual private cloud (VPC), allowing you to control your network settings. If you select this option, you must specify an existing VPC endpoint by choosing its VPC endpoint ID from the dropdown list. You can also provide its VPC endpoint Domain Name System (DNS) name or IP address.

3. In the **Gateway connection options** section, for **Connection options**, choose how to identify your gateway to AWS. You can choose from the following options:

   - **IP address** – Provide the IP address of your gateway in the corresponding field. This IP address must be public or accessible from within your current network, and you must be able to connect to it from your web browser.

     You can obtain the gateway IP address by logging into the gateway's local console from your hypervisor client, or by copying it from your Amazon EC2 instance details page.

   - **Activation key** – Provide the activation key for your gateway in the corresponding field. You can generate an activation key using the gateway's local console. If your gateway's IP address is unavailable, choose this option.

4. Now that you have chosen how you want your gateway to connect to AWS, you must activate the gateway. Choose **Next** to proceed.

Review settings and activate your Amazon FSx File Gateway

To activate a new FSx File Gateway

1. If you have not done so already, complete the procedures described in the following topics:

   - Set up an Amazon FSx File Gateway
   - Connect your Amazon FSx File Gateway to AWS

   When finished, choose **Next** to open the **Review and activate** page in the AWS Storage Gateway console.

2. Review the initial gateway details for each section on the page.

3. If a section contains errors, choose **Edit** to return to the corresponding settings page and make changes.

> **Important**
>
> You cannot modify the gateway options or connection settings after your gateway is activated.

4. Now that you have activated your gateway, you must perform the first-time configuration to allocate local storage disks and configure logging. Choose **Next** to proceed.

Configure your Amazon FSx File Gateway

To perform the first-time configuration on a new FSx File Gateway

1. If you have not done so already, complete the procedures described in the following topics:

   - Set up an Amazon FSx File Gateway
   - Connect your Amazon FSx File Gateway to AWS
   - Review settings and activate your Amazon FSx File Gateway

   When finished, choose **Next** to open the **Configure gateway** page in the AWS Storage Gateway console.

2. In the **Configure storage** section, use the dropdown lists to allocate at least one local disk with at least 150 gibibytes (GiB) capacity to **Cache**. The local disks listed in this section correspond to the physical storage that you provisioned on your host platform.

3. In the **CloudWatch log group** section, choose how to set up Amazon CloudWatch Logs to monitor the health of your gateway. You can choose from the following options:

   - **Create a new log group** – Set up a new log group to monitor your gateway.
   - **Use an existing log group** – Choose an existing log group from the corresponding dropdown list.
   - **Deactivate logging** – Do not use Amazon CloudWatch Logs to monitor your gateway.

> **Note**
>
> To receive Storage Gateway health logs, the following permissions must be present in your log group resource policy. Replace the *highlighted section* with the specific log group resourceArn information for your deployment.
>
> ```
> "Sid": "AWSLogDeliveryWrite20150319",
>       "Effect": "Allow",
>       "Principal": {
>         "Service": [
>           "delivery.logs.amazonaws.com"
>         ]
>       },
>       "Action": [
>         "logs:CreateLogStream",
>         "logs:PutLogEvents"
>       ],
>       "Resource": "arn:aws:logs:eu-west-1:1234567890:log-group:/foo/bar:log-stream:*"
> ```
>
> The "Resource" element is required only if you want the permissions to apply explicitly to an individual log group.

4. In the **CloudWatch alarms** section, choose how to set up Amazon CloudWatch alarms to notify you when your gateway's metrics deviate from defined limits. You can choose from the following options:

   - **Create Storage Gateway's recommended alarms** – Create all recommended CloudWatch alarms automatically when the gateway is created. For more information on recommended alarms, see Understanding CloudWatch alarms.

     > **Note**
     >
     > This feature requires CloudWatch policy permissions which are *not* automatically granted as part of the preconfigured Storage Gateway full access policy. Make sure your security policy grants the following permissions before you attempt to create recommended CloudWatch alarms:
     >
     > - `cloudwatch:PutMetricAlarm` - create alarms

- `cloudwatch:DisableAlarmActions` - turn alarm actions off
- `cloudwatch:EnableAlarmActions` - turn alarm actions on
- `cloudwatch:DeleteAlarms` - delete alarms

- **Create a custom alarm** – Configure a new CloudWatch alarm to be notified about your gateway's metrics. Choose **Create alarm** to define metrics and specify alarm actions in the Amazon CloudWatch console. For instructions, see Using Amazon CloudWatch alarms in the *Amazon CloudWatch User Guide*.
- **No alarm** – Do not use CloudWatch alarms to be notified about your gateway's metrics.

5. (Optional) In the **Tags** section, choose **Add new tag**, then enter a case-sensitive key-value pair to help you search and filter for your gateway on the list pages in the AWS Storage Gateway console. Repeat this step to add as many tags as you need.
6. (Optional) In the **Verify VMware High Availability configuration** section, if your gateway is deployed on a VMware host that is part of a VMware High Availability (HA) cluster, choose **Verify VMware HA** to test whether the HA configuration is working properly.

   Note

   This section appears only for gateways that are running on the VMware host platform. This step is not required to complete the gateway configuration process. You can test your gateway's HA configuration at any time. Verification takes a few minutes, and reboots the Storage Gateway virtual machine (VM).

7. Choose **Configure** to finish creating your gateway.

   To check the status of your new gateway, search for it on the **Gateway overview** page of the AWS Storage Gateway console.

Now that you have created your gateway, you must attach a file system for it to use. For instructions, see Attach an Amazon FSx for Windows File Server file system.

If you do not have an existing Amazon FSx file system to attach, you must create one. For instructions, see Getting started with Amazon FSx.

Activating a gateway in a virtual private cloud

You can create a private connection between your on-premises gateway appliance and cloud-based storage infrastructure. You can use this connection to activate your gateway and configure it to transfer data to AWS storage services without communicating over the public internet. Using the Amazon VPC service, you can launch AWS resources, including private network interface endpoints, in a custom virtual private cloud (VPC). A VPC gives you control over network settings such as IP address range, subnets, route tables, and network gateways. For more information about VPCs, see What is Amazon VPC? in the *Amazon VPC User Guide*.

To activate your gateway in a VPC, use the Amazon VPC Console to create a VPC endpoint for Storage Gateway and get the VPC endpoint ID, then specify this VPC endpoint ID when you create and activate the gateway. For more information, see Connect your Amazon FSx File Gateway to AWS.

To configure your FSx File Gateway to transfer data through the VPC, you must establish a VPN or AWS DirectConnect link between the Amazon FSx for Windows File Server VPC and the network where your gateway is deployed.

> **Note**
>
> You must activate your gateway in the same region where you create the VPC endpoint for Storage Gateway.

Create a VPC endpoint for Storage Gateway

Follow these instructions to create a VPC endpoint. If you already have a VPC endpoint for Storage Gateway, you can use it.

To create a VPC endpoint for Storage Gateway

1. Sign in to the AWS Management Console and open the Amazon VPC console at https://console.aws.amazon.com/vpc/.
2. In the navigation pane, choose **Endpoints**, and then choose **Create Endpoint**.
3. On the **Create Endpoint** page, choose **AWS Services** for **Service category**.
4. For **Service Name**, choose `com.amazonaws.`*`region`*`.storagegateway`. For example `com.amazonaws.us-east-2.storagegateway`.

5. For **VPC**, choose your VPC and note its Availability Zones and subnets.
6. Verify that **Enable Private DNS Name** is not selected.
7. For **Security group**, choose the security group that you want to use for your VPC. You can accept the default security group. Verify that all of the following TCP ports are allowed in your security group:

   - TCP 443
   - TCP 1026
   - TCP 1027
   - TCP 1028
   - TCP 1031
   - TCP 2222
8. Choose **Create endpoint**. The initial state of the endpoint is **pending**. When the endpoint is created, note the ID of the VPC endpoint that you just created.
9. When the endpoint is created, choose **Endpoints**, then choose the new VPC endpoint.
10. In **Details** tab of the selected storage gateway endpoint, under **DNS Names**, use the first DNS name that doesn't specify an Availability Zone. Your DNS name should look similar to the following example: `vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com`

Now that you have a VPC endpoint, you can create and activate your gateway. For more information, see Create and activate an Amazon FSx File Gateway.

For information about getting an activation key, see Getting an activation key for your gateway.

Configure Microsoft Active Directory domain access settings

In this step, you configure access settings to join your Amazon FSx File Gateway to a Microsoft Active Directory domain.

To configure Active Directory settings

1. In the Storage Gateway console, choose **FSx file systems** from the navigation menu.
2. Choose **Attach FSx file system**.
3. On the **Confirm gateway** page, choose the gateway you want to join to your Active Directory domain from the drop-down menu.

   If you don't have a gateway, you must create one. Make sure your gateway can resolve the name of your Active Directory Domain Controller. For information, see Prerequisites.
4. Enter values for the **Active Directory settings**:

   > **Note**
   >
   > If your gateway is already joined to a domain, you don't need to join again. Go to the next step.

   - For **Domain name**, enter the domain name of the Active Directory that you want to use.
   - For **Domain user**, enter the user name of the Active Directory user that you want to use to join the gateway to the domain. This user must have the necessary permissions. For more information, see Active Directory service account permission requirements.
   - For **Domain password**, enter the password for the user.
   - For **Organizational unit- optional**, you can specify an organizational unit the Active Directory belongs to.

     > **Note**
     >
     > If you leave this field blank, joining a domain creates an Active Directory computer account in the default computers container (which is not an OU), using the gateway's

Gateway ID as the account name (for example, SGW-1234ADE). It is not possible to customize the name of this account.
If your Active Directory environment requires that you pre-stage accounts to facilitate the join domain process, you will need to create this account ahead of time.
If your Active Directory environment has a designated OU for new computer objects, you must specify that OU when joining the domain.

- Enter a value for **Domain controller(s) - optional**.

5. Choose **Next** to open the **Attach FSx File system** page.

Next step

Attach an Amazon FSx for Windows File Server file system

Attach an Amazon FSx for Windows File Server file system

You must have an FSx for Windows File Server file system before you can attach it to an FSx File Gateway. If you don't have a file system, you must create one. For instructions, see Step 1: Create Your File System in the *Amazon FSx for Windows File Server User Guide*.

The next step is to attach an Amazon FSx file system to the gateway. When you attach an Amazon FSx file system, all the file shares on the file system are made available to Amazon FSx File Gateway (FSx File Gateway) for you to mount.

Note

The following limitations apply when writing to an Amazon FSx file system from Amazon FSx File Gateway:

- Your Amazon FSx file system and your FSx File Gateway must be owned by the same AWS account and located in the same AWS Region.
- Each gateway can support up to five attached file systems. When you're attaching a file system, the Storage Gateway console notifies you if the selected gateway is at capacity. In that case, you must choose a different gateway or detach a file system before you can attach another one.
- FSx File Gateway supports soft storage quotas (which warn you when users surpass their data limits), but does not support hard quotas (which enforce data limits by denying write access). Soft quotas are supported for all users except the Amazon FSx admin user. For more information about setting up storage quotas, see Storage quotas in the Amazon FSx User Guide.
- We don't recommend using Microsoft Distributed File System (DFS) to redirect users to your Amazon FSx file system through FSx File Gateway. Instead, configure DFS to redirect directly to the Amazon FSx file system in the AWS Cloud as described in Grouping multiple file systems with DFS Namespaces in the *Amazon FSx for Windows File Server User Guide*.

To attach an Amazon FSx file system

1. In the Storage Gateway console, on the **FSx file systems** > **Attach FSx file system** page, complete the following fields in the **FSx file system settings** section:

   - For **FSx file system name**, choose the file system that you want to attach from the dropdown list.
   - For **Local Endpoint IP address**, enter the gateway IP address that clients will use to browse file shares on the FSx file system.

   > **Note**
   >
   > - You must specify an IP address for each file system attached to the gateway.
   > - For Amazon EC2 gateways, you can specify the private IP address of the EC2 instance, unless it is already in use by a different file system, in which case you must add a new private address to the gateway, then restart it. For more information, see Multiple IP addresses in the *Amazon EC2 User Guide*.
   > - For on-premises gateways, you can specify the IP address of the primary network interface (static or DHCP), unless it is already in use by a different file system, in which case you must provide a different IP address from the same subnet as the primary interface, which will be made available as a virtual IP. Do not use an IP address assigned to any network interface other than the primary.

2. In the **Service account settings** section, provide the service account sign-in credentials that is associated with the Amazon FSx file system.

   > **Note**
   >
   > This service account must have Backup Operators privileges from the Active Directory service that is associated with your Amazon FSx file systems or have equivalent permissions.

   > **Important**
   >
   > To ensure sufficient permissions to files, folders, and file metadata, we recommend that you make the service account a member of the file system administrators group.

If you are using AWS Directory Service for Microsoft Active Directory with Amazon FSx for Windows File Server, the service account must be a member of the AWS Delegated FSx Administrators group.
If you are using a self-managed Active Directory with Amazon FSx for Windows File Server, we recommend that the service account be a member of the *custom delegated file system administrators* group you specified for file system administration when you created your Amazon FSx file system.
If you chose not to create a *custom delegated file system administrators* group when you created the Amazon FSx filesystem, the default group is *Domain Admins*. While you can make the service account a member of this group instead, it is not recommended as a best practice.
For more information, see Delegating privileges to your Amazon FSx service account in the *Amazon FSx for Windows File Server User Guide*.

3. In the **Audit logs** section, choose **Existing log groups**, and choose the log that you want to use to monitor access to your Amazon FSx file system. You can create a new one. If you don't want to monitor your system, choose **Disable logging**.
4. For **Automated cache refresh setting**, if you want your cache to refresh automatically, choose **Set refresh interval** and specify an interval between 5 minutes and 30 days.
5. (Optional) In the **Tags** section, choose **Add new tag** to add one or more keys and a value for tagging your settings.
6. Choose **Next** and review the settings. To change your settings, you can choose **Edit** in each section.
7. When you are done, choose **Finish**.

Next step

Mount and use your Amazon FSx file share

Mount and use your Amazon FSx file share

Before mounting your file share on the client, wait for the status of the Amazon FSx file system to change to **Available**. After your file share is mounted, you can start using your Amazon FSx File Gateway (FSx File Gateway).

Topics

Mount your SMB file share on your client

In this step, you mount your SMB file share and map to a drive that is accessible to your client. The console's File Gateway section shows the supported mount commands that you can use for SMB clients. Following are some additional options to try.

You can use several different methods for mounting SMB file shares, including the following:

- The `net use` command – Doesn't persist across system reboots, unless you use the `/persistent:(yes:no)` switch.
- The `CmdKey` command line utility – Creates a persistent connection to a mounted SMB file share that remains after a reboot.
- A network drive mapped in File Explorer – Configures the mounted file share to reconnect at sign-in and to require that you enter your network credentials.
- PowerShell script – Can be persistent, and can be either visible or invisible to the operating system while mounted.

> **Note**
>
> If you are a Microsoft Active Directory user, check with your administrator to ensure that you have access to the SMB file share before mounting the file share to your local system. Amazon FSx File Gateway doesn't support SMB locking or SMB extended attributes.

To mount an SMB file share for Active Directory users using the net use command

1. Make sure that you have access to the SMB file share before mounting the file share to your local system.
2. For Microsoft Active Directory clients, enter the following command at the command prompt:

   ```
   net use [WindowsDriveLetter]: \\[Gateway IP Address]\[Name of the share on the FSx file system]
   ```

To mount an SMB file share on Windows using CmdKey

1. Press the Windows key and enter **cmd** to view the command prompt menu item.
2. Open the context (right-click) menu for **Command Prompt**, and choose **Run as administrator**.
3. Enter the following command:

   ```
   C:\>cmdkey /add:[Gateway VM IP address] /user:[DomainName]\[UserName] /pass:[Password]
   ```

> **Note**
>
> When mounting file shares, you might need to remount your file share after rebooting your client.

To mount an SMB file share using Windows File Explorer

1. Press the Windows key and enter **File Explorer** in the **Search Windows** box, or press **Win +E**.
2. In the navigation pane, choose **This PC**. Then, on the **Computer** tab, choose **Map network drive**.
3. In the **Map network drive** dialog box, choose a drive letter for **Drive**.
4. For **Folder**, enter *****[File Gateway IP]****[SMB File Share Name]**, or choose **Browse** to select your SMB file share from the dialog box.
5. (Optional) Select **Reconnect at sign-up** if you want your mount point to persist after reboots.
6. (Optional) Select **Connect using different credentials** if you want a user to enter the Active Directory logon or guest account user password.

7. Choose **Finish** to complete your mount point.

Test your FSx File Gateway

You can copy files and directories to your mapped drive. The files automatically upload to your FSx for Windows File Server file system.

To upload files from your Windows client to Amazon FSx

1. On your Windows client, navigate to the drive that you mounted your file system on. The name of the drive is preceded by the name of your file system.
2. Copy files or a directory to the drive.

> **Note**
>
> File Gateways don't support creating hard or symbolic links on a file share.

Managing your Amazon FSx File Gateway resources

The following sections provide information about how to manage your Amazon FSx File Gateway (FSx File Gateway) resources, including attaching and detaching Amazon FSx file systems, and configuring Microsoft Active Directory settings.

Topics

Understanding gateway status

Each gateway in your AWS Storage Gateway deployment has an associated status that tells you at a glance what the health of the gateway is. Most of the time, the status indicates that the gateway is functioning normally and that no action is needed on your part. In some cases, the status indicates a problem that might or might not require action on your part.

You can see the status for each gateway in your deployment on the **Gateways** page of the Storage Gateway console. The gateway status appears in the **Status** column next to the name of the gateway. A gateway that is functioning normally has a status of `RUNNING`.

In the following table, you can find a description of each gateway status, and whether you should act based on the status. A gateway should have `RUNNING` status all or most of the time it's in use.

Status	Meaning
`RUNNING`	The gateway is configured properly and is available to use.
`OFFLINE`	Your gateway might be in an `OFFLINE` status for one or more of the following reasons:

Status	Meaning
	• The gateway can't reach the Storage Gateway service endpoints. • The gateway had an unexpected shutdown. • The gateway has an associated cache disk that is disconnected, has been modified, or has failed.

Understanding file system status

You can view the health of a file system at a glance by looking at its status. If the status indicates that the file system is functioning normally, no action is needed on your part. If the status indicates that there's a problem, you can investigate to determine whether action could be required.

You can view a file system's status on the Storage Gateway console in the **Status** column. A file system that's functioning properly shows a status of AVAILABLE. This should be the status most of the time.

The following table describes file share statuses, what they mean, and whether action might be required.

Status	Meaning
AVAILABLE	The file system is configured properly and is available to use. This is the standard status for a file system that's working properly.
CREATING	The file system is not yet fully created and is not ready for use. The CREATING status is transitional. No action is required. If the file system gets stuck in this status, it's probably because the gateway VM lost connection to AWS.
UPDATING	The file system configuration is currently updating. The UPDATING status is transitional. No action is required. If a file system gets stuck in this status, it's probably because the gateway VM lost connection to AWS.

Status	Meaning
DELETING	The file system is being deleted. The file system is not deleted until all data is uploaded to AWS. The DELETING status is transitional, and no action is required.
FORCE_DELETING	The file system is being deleted forcibly. The file system is deleted immediately and data is not uploaded to AWS. The FORCE_DELETING status is transitional, and no action is required.
ERROR	The file system is in an unhealthy state. Action is required. Some possible causes include problems with access credentials or privileges, connectivity issues, or insufficient storage space on the file system. When the issue that caused the unhealthy state is resolved, the file system returns to a status of AVAILABLE.

Edit basic information for an FSx File Gateway

You can use the Storage Gateway console to edit basic information for an existing gateway, including the gateway name, time zone, and CloudWatch log group.

To edit basic information for an existing gateway

1. Open the Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. Choose **Gateways**, then choose the gateway for which you want to edit basic information.
3. From the **Actions** dropdown menu, choose **Edit gateway information**.
4. For **Gateway name**, enter a name for your gateway. You can search for this name to find your gateway on the list pages in the Storage Gateway console.

> **Note**
>
> Gateway names must be between 2 and 255 characters, and cannot include a slash (\ or /).

Changing a gateway's name will disconnect any CloudWatch alarms set up to monitor the gateway. To reconnect the alarms, update the **GatewayName** for each alarm in the CloudWatch console.

5. For **Gateway time zone**, choose the local time zone for the part of the world where you want to deploy your gateway.
6. For **Choose how to set up log group**, choose how to set up Amazon CloudWatch Logs to monitor the health of your gateway. You can choose from the following options:
   - **Create a new log group** – Set up a new log group to monitor your gateway.
   - **Use an existing log group** – Choose an existing log group from the corresponding dropdown list.
   - **Deactivate logging** – Do not use Amazon CloudWatch Logs to monitor your gateway.
7. When you finish modifying the settings you want to change, choose **Save changes**.

Set a security level for your gateway

You can configure the SMB security level for your FSx File Gateway to specify whether the gateway should require Server Message Block (SMB) signing or SMB encryption.

To configure security level

1. Open the Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. Choose **Gateways**, then choose the gateway for which you want to edit SMB settings.
3. From the **Actions** dropdown menu, choose **Edit SMB settings**, then choose **SMB security settings**.
4. For **Security level**, choose one of the following:

Note

For information about configuring this setting using the AWS API, see UpdateSMBSecurityStrategy in the *AWS Storage Gateway API Reference*.
A higher security level can affect performance of the gateway.

- **Mandatory encryption** – If you choose this option, FSx File Gateway only allows connections from SMBv3 clients that use 256-bit AES encryption algorithms. 128-bit algorithms are not allowed. This option is recommended for environments that handle sensitive data. It works with SMB clients on Microsoft Windows 8, Windows Server 2012, or later.
- **Enforce encryption** – If you choose this option, FSx File Gateway only allows connections from SMBv3 clients that have encryption turned on. Both 256-bit and 128-bit algorithms are allowed. This option is recommended for environments that handle sensitive data. It works with SMB clients on Microsoft Windows 8, Windows Server 2012, or later.
- **Enforce signing** – If you choose this option, FSx File Gateway only allows connections from SMBv2 or SMBv3 clients that have signing turned on. This option works with SMB clients on Microsoft Windows Vista, Windows Server 2008, or later.

> **Note**
>
> The default security level for FSx File Gateway is **Enforce encryption**.

5. Choose **Save**.

Editing Active Directory settings for n FSx File Gateway

To use your corporate Microsoft Active Directory or AWS Managed Microsoft AD for user authenticated access to your Amazon FSx file system, edit the SMB settings for your gateway and provide your Active Directory domain credentials. Doing this allows your gateway to join your Active Directory domain and allows members of the domain to access the file system.

> **Note**
>
> Using AWS Directory Service, you can create a hosted Active Directory domain service in the AWS Cloud.
> To use AWS Managed Microsoft AD with an Amazon EC2 gateway, you must create the Amazon EC2 instance in the same VPC as the AWS Managed Microsoft AD, add the _workspaceMembers security group to the Amazon EC2 instance, and join the AD domain using the Admin credentials from the AWS Managed Microsoft AD.
> For more information about AWS Managed Microsoft AD, see the *AWS Directory Service Administration Guide*.

For more information about Amazon EC2, see the *Amazon Elastic Compute Cloud Documentation*.

To turn on Active Directory authentication

1. Open the Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. Choose **Gateways**, then choose the gateway for which you want to edit SMB settings.
3. From the **Actions** drop-down menu, choose **Edit SMB settings**, then choose **Active Directory settings**.
4. For **Domain name**, enter the name of the Active Directory domain you want your gateway to join.

   Note

   Active Directory status shows **Detached** when a gateway has never joined a domain.
   Your Active Directory service account must have the requisite permissions. For more information, see Active Directory service account permission requirements.
   Joining a domain creates an Active Directory computer account in the default computers container (which is not an OU), using the gateway's **Gateway ID** as the account name (for example, SGW-1234ADE). It is not possible to customize the name of this account.
   If your Active Directory environment requires that you pre-stage accounts to facilitate the join domain process, you will need to create this account ahead of time.
   If your Active Directory environment has a designated OU for new computer objects, you must specify that OU when joining the domain.
   If your gateway can't join an Active Directory directory, try joining with the directory's IP address by using the JoinDomain API operation.
5. For **Domain user** and **Domain password**, enter the credentials for the Active Directory service account that the gateway will use to join the domain.
6. (Optional) For **Organization unit (OU)**, enter the designated OU that your Active Directory uses for new computer objects.

7. (Optional) For **Domain controller(s) (DC)**, enter the name of one or more DCs through which your gateway will connect to Active Directory. You can enter multiple DCs as a comma-separated list. You can leave this field blank to allow DNS to automatically select a DC.
8. Choose **Save changes**.

Editing settings for an Amazon FSx file system

After creating an Amazon FSx for Windows File Server file system, you can edit settings for CloudWatch logs, automated cache refresh, and Amazon FSx service account credentials.

To edit Amazon FSx file system settings

1. Open the Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. In the navigation pane, choose **File system**, and choose the file system whose settings you want to edit.
3. For **Actions**, choose **Edit file system settings**.
4. In the file system settings section, verify the gateway, Amazon FSx location, and IP address information.

   > **Note**
   >
   > You cannot edit a file system's IP address after it is attached to a gateway. To change the IP address, you must detach and reattach the file system.

5. In the **Audit logs** section, choose an option to use CloudWatch log groups to monitor access to Amazon FSx file systems. You can use an existing log group.
6. For **Automated cache refresh settings**, choose an option. If you choose **Set refresh interval**, set the time in days, hours, and minutes to refresh the file system's cache using Time To Live (TTL).

   TTL is the length of time since the last refresh. When the directory is accessed after that length of time, the File Gateway refreshes that directory's contents from the Amazon FSx file system.

> **Note**
>
> Valid refresh interval values are between 5 minutes and 30 days.

7. In the **Service account settings - optional** section, enter a user name and a **Password**. These credentials are for a user that has the Backup Administrator role from the Active Directory service associated with your Amazon FSx file systems.
8. Choose **Save changes**.

Detaching an Amazon FSx file system

Detaching a file system doesn't delete your data in FSx for Windows File Server. Data that is written to these the file systems before you detach them will still be uploaded to your FSx for Windows File Server.

To detach an Amazon FSx file system

1. Open the Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. Choose **FSx file systems**, then select one or more file systems to detach.
3. For **Actions**, choose **Detach file system**. The confirmation dialog box appears.
4. Verify that you want to detach the specified file systems, then type the word *detach* in the confirmation box and choose **Detach**.

Monitoring Storage Gateway

The topics in this section describe how to monitor a gateway using Amazon CloudWatch, including monitoring cache storage and other resources associated with the gateway. You use the Storage Gateway console to view metrics and alarms for your gateway. For example, you can view the number of bytes used in read and write operations, the time spent in read and write operations, and the time taken to retrieve data from the AWS Cloud. With metrics, you can track the health of your gateway and set up alarms to notify you when one or more metrics fall outside a defined threshold.

Storage Gateway provides CloudWatch metrics at no additional charge. Storage Gateway metrics are recorded for a period of two weeks. By using these metrics, you can access historical information and get a better perspective on how your gateways are performing. Storage Gateway also provides CloudWatch alarms, except high-resolution alarms, at no additional charge. For more information about CloudWatch pricing, see Amazon CloudWatch pricing. For more information about CloudWatch, see the Amazon CloudWatch User Guide.

Topics

- Understanding CloudWatch alarms - Learn basic information about CloudWatch alarms, including alarm states and recommended configurations.
- Create recommended CloudWatch alarms - Learn how you can quickly and automatically configure all recommended CloudWatch alarms as part of the initial File Gateway setup process.
- Create a custom CloudWatch alarm - Learn how you can create a custom CloudWatch alarm to monitor a specific metric using specific evaluation criteria to trigger alarm states and send notifications.
- Monitoring your FSx File Gateway - Learn how to view CloudWatch logs and audit logs, and find information about the specific gateway and file sharefile system metrics that are reported by your gateway.

Understanding CloudWatch alarms

CloudWatch alarms monitor information about your gateway based on metrics and expressions. You can add CloudWatch alarms for your gateway and view their statuses in the Storage Gateway console. For more information about the metrics that are used to monitor FSx File Gateway, see

Understanding gateway metrics and Understanding file system metrics. For each alarm, you specify conditions that will activate its ALARM state. Alarm status indicators in the Storage Gateway console turn red when in the ALARM state, making it easier for you to monitor status proactively. You can configure alarms to invoke actions automatically based on sustained changes in state. For more information about CloudWatch alarms, see Using Amazon CloudWatch alarms in the *Amazon CloudWatch User Guide*.

> **Note**
>
> If you don't have permission to view CloudWatch, you can't view the alarms.

For each activated gateway, we recommend that you create the following CloudWatch alarms:

- High IO wait: `IoWaitpercent` >= 20 for 3 datapoints in 15 minutes
- Cache percent dirty: `CachePercentDirty` > 80 for 4 datapoints within 20 minutes
- Files failing upload: `FilesFailingUpload` >= 1 for 1 datapoint within 5 minutes
- File system error: `FileSystem-ERROR` >= 1 for 1 datapoint within 5 minutes
- Health notifications: `HealthNotifications` >= 1 for 1 datapoints within 5 minutes. When configuring this alarm, set **Missing data treatment** to **notBreaching**.

  > **Note**
  >
  > You can set a health notification alarm only if the gateway had a previous health notification in CloudWatch.

For gateways on VMware host platforms that are part of a VMware High Availability cluster, we also recommend this additional CloudWatch alarm:

- Availability notifications: `AvailabilityNotifications` >= 1 for 1 datapoints within 5 minutes. When configuring this alarm, set **Missing data treatment** to **notBreaching**.

The following table describes CloudWatch alarm states.

State	Description
OK	The metric or expression is within the defined threshold.
Alarm	The metric or expression is outside of the defined threshold.
Insufficient data	The alarm has just started, the metric is not available, or not enough data is available for the metric to determine the alarm state.
None	No alarms are created for the gateway. To create a new alarm, see Create a custom CloudWatch alarm for your gateway.
Unavailable	The state of the alarm is unknown. Choose **Unavailable** to view error information in the **Monitoring** tab.

Creating recommended CloudWatch alarms for your gateway

When you create a new gateway using the Storage Gateway console, you can choose to create all recommended CloudWatch alarms automatically as part of the initial setup process. For more information, see Configure your Amazon FSx File Gateway. If you want to add or update recommended CloudWatch alarms for an existing gateway after you have already completed the first-time setup, use the following procedure.

To add or update recommended CloudWatch alarms for an existing gateway

> **Note**
>
> This feature requires CloudWatch policy permissions, which are *not* automatically granted as part of the preconfigured Storage Gateway full access policy. Make sure your security policy grants the following permissions before you attempt to create recommended CloudWatch alarms:
>
> - `cloudwatch:PutMetricAlarm` - create alarms

- `cloudwatch:DisableAlarmActions` - turn alarm actions off
- `cloudwatch:EnableAlarmActions` - turn alarm actions on
- `cloudwatch:DeleteAlarms` - delete alarms

1. Open the Storage Gateway console at https://console.aws.amazon.com/storagegateway/home/.
2. In the navigation pane on the left side of the page, choose **Gateways**, and then choose the gateway for which you want to create recommended CloudWatch alarms.
3. On the **Details** page for the gateway, choose the **Monitoring** tab.
4. Under **Alarms**, choose **Create recommended alarms**. The recommended alarms are created automatically.

   The **Alarms** section lists all CloudWatch alarms for a specific gateway. From here, you can select and delete one or more alarms, turn alarm actions on or off, and create new alarms.

Create a custom CloudWatch alarm for your gateway

CloudWatch uses Amazon Simple Notification Service (Amazon SNS) to send alarm notifications when an alarm changes state. An alarm watches a single metric over a time period that you specify, and 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 that's sent to an Amazon SNS topic. You can create an Amazon SNS topic when you create a CloudWatch alarm. For more information about Amazon SNS, see What is Amazon SNS? in the *Amazon Simple Notification Service Developer Guide*.

To create a CloudWatch alarm in the Storage Gateway console

1. Open the Storage Gateway console at https://console.aws.amazon.com/storagegateway/home/.
2. In the navigation pane, choose **Gateways**, then choose the gateway for which you want to create an alarm.
3. On the gateway details page, choose the **Monitoring** tab.
4. Under **Alarms**, choose **Create alarm** to open the CloudWatch console.
5. Use the CloudWatch console to create the type of alarm that you want. You can create the following types of alarms:

- Static threshold alarm: An alarm based on a set threshold for a chosen metric. The alarm enter the ALARM state when the metric breaches the threshold for a specified number of evaluation periods.

  To create a static threshold alarm, see Creating a CloudWatch alarm based on a static threshold in the *Amazon CloudWatch User Guide*.

- Anomaly detection alarm: Anomaly detection mines past metric data and creates a model of expected values. You set a value for the anomaly detection threshold, and CloudWatch uses this threshold with the model to determine the "normal" range of values for the metric. A higher value for the threshold produces a thicker band of "normal" values. You can choose to activate the alarm only when the metric value is above the band of expected values, only when it's below the band, or when it's above or below the band.

  To create an anomaly detection alarm, see Creating a CloudWatch alarm based on anomaly detection in the *Amazon CloudWatch User Guide*.

- Metric math expression alarm: An alarm based one or more metrics used in a math expression. You specify the expression, threshold, and evaluation periods.

  To create a metric math expression alarm, see Creating a CloudWatch alarm based on a metric math expression in the *Amazon CloudWatch User Guide*.

- Composite alarm: An alarm that determines its alarm state by watching the alarm states of other alarms. A composite alarm can help you reduce alarm noise.

  To create a composite alarm, see Creating a composite alarm in the *Amazon CloudWatch User Guide*.

6. After you create the alarm in the CloudWatch console, return to the Storage Gateway console. You can view the alarm by doing one of the following:

   - In the navigation pane, choose **Gateways**, then choose the gateway for which you want to view alarms. On the **Details** tab, under **Alarms**, choose **CloudWatch Alarms**.
   - In the navigation pane, choose **Gateways**, choose the gateway for which you want to view alarms, then choose the **Monitoring** tab.

     The **Alarms** section lists all of the CloudWatch alarms for a specific gateway. From here, you can select and delete one or more alarms, turn alarm actions on or off, and create new alarms.

- In the navigation pane, choose **Gateways**, then choose the alarm state of the gateway for which you want to view alarms.

For information about how to edit or delete an alarm, see Editing or deleting a CloudWatch alarm.

> **Note**
>
> When you delete a gateway using the Storage Gateway console, all CloudWatch alarms associated with the gateway are also automatically deleted.

Monitoring your FSx File Gateway

You can monitor your FSx File Gateway and associated resources in AWS Storage Gateway by using Amazon CloudWatch metrics and audit logs. You can also use CloudWatch Events to get notified when your file operations are done.

Topics

Getting FSx File Gateway health logs with CloudWatch log groups

You can use Amazon CloudWatch Logs to get information about the health of your FSx File Gateway and related resources. You can use the logs to monitor your gateway for errors that it encounters. In addition, you can use Amazon CloudWatch subscription filters to automate processing of the log information in real time. For more information, see Real-time Processing of Log Data with Subscriptions in the *Amazon CloudWatch User Guide*.

For example, you can configure a CloudWatch log group to monitor your gateway and get notified when your FSx File Gateway fails to upload files to an Amazon FSx file system. You can configure the group either when you are activating the gateway or after your gateway is activated and up and running. For information about how to configure a CloudWatch log group when activating a

gateway, see Configure your Amazon FSx File Gateway. For general information about CloudWatch log groups, see Working with Log Groups and Log Streams in the *Amazon CloudWatch User Guide.*

For information about how to troubleshoot the errors that may be reported by FSx File Gateway, see Troubleshooting: File Gateway issues.

Configuring a CloudWatch log group after your gateway is activated

The following procedure shows you how to configure a CloudWatch Log Group after your gateway is activated.

To configure a CloudWatch log group to work with your FSx File Gateway

1. Sign in to the AWS Management Console and open the Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. In the navigation pane, choose **Gateways**, and then choose the gateway that you want to configure the CloudWatch log group for.
3. For **Actions**, choose **Edit gateway information**.
4. For **Choose how to set up log group**, choose one of the following:

   - **Create a new log group** to create a new CloudWatch log group.
   - **Use an existing log group** to use a CloudWatch log group that already exists.

     Choose a log group from the **Existing log group list**.
   - **Deactivate logging** if you don't want to monitor your gateway using CloudWatch log groups.
5. Choose **Save changes**.
6. To see the health logs for your gateway, do the following:

   1. In the navigation pane, choose **Gateways**, and then choose the gateway that you configured the CloudWatch log group for.
   2. Choose the **Details** tab, and under **Health logs**, choose **CloudWatch Logs**. The **Log group details** page opens in the CloudWatch console.

Using Amazon CloudWatch metrics

You can get monitoring data for your FSx File Gateway by using either the AWS Management Console or the CloudWatch API. The console displays a series of graphs based on the raw data

from the CloudWatch API. The CloudWatch API can also be used through one of the AWS SDKs or Amazon CloudWatch API tools. Depending on your needs, you might prefer to use either the graphs displayed in the console or retrieved from the API.

Regardless of which method you use to work with metrics, you must specify the following information:

- The metric dimension to work with. A *dimension* is a name-value pair that helps you to uniquely identify a metric. The dimensions for Storage Gateway are `GatewayId` and `GatewayName`. In the CloudWatch console, you can use the `Gateway Metrics` view to select gateway-specific dimensions. For more information about dimensions, see Dimensions in the *Amazon CloudWatch User Guide*.
- The metric name, such as `ReadBytes`.

The following table summarizes the types of Storage Gateway metric data that are available to you.

Amazon CloudWatch namespace	Dimension	Description
`AWS/StorageGateway`	`GatewayId`, `GatewayName`	These dimensions filter for metric data that describes aspects of the gateway. You can identify a FSx File Gateway to work with by specifying both the `GatewayId` and the `GatewayName` dimensions. Throughput and latency data of a gateway are based on all the file shares in the gateway. Data is available automatically in 5-minute periods at no charge.

Working with gateway and file metrics is similar to working with other service metrics. You can find a discussion of some of the most common metrics tasks in the CloudWatch documentation listed following:

- Viewing available metrics

- Getting statistics for a metric
- Creating CloudWatch alarms

Understanding gateway metrics

The following table describes metrics that cover FSx File Gateways. Each gateway has a set of metrics associated with it. Some gateway-specific metrics have the same name as certain file-system-specific metrics. These metrics represent the same kinds of measurements, but are scoped to the gateway rather than the file system.

Always specify whether you want to work with a gateway or a file system when working with a particular metric. Specifically, when working with gateway metrics, you must specify the `Gateway Name` for the gateway whose metric data you want to view. For more information, see Using Amazon CloudWatch metrics.

> **Note**
>
> Some metrics return data points only when new data has been generated during the most recent monitoring period.

The following table describes the metrics that you can use to get information about your FSx File Gateways.

Metric	Description
`AvailabilityNotifications`	This metric reports the number of availabil ity-related health notifications that were generated by the gateway in the reporting period. Units: Count
`CacheDirectorySize`	This metric tracks the size of folders in the gateway cache. Folder size is determined by the number of files and subfolders in its first level, this does not count recursively into subfolders.

Metric	Description
	Use this metric with the `Average` statistic to measure the average size of a folder in the gateway cache. Use this metric with the `Max` statistic to measure the maximum size of a folder in the gateway cache. Units: Count
`CacheFileSize`	This metric tracks the size of files in the gateway cache. Use this metric with the `Average` statistic to measure the average size of a file in the gateway cache. Use this metric with the `Max` statistic to measure the maximum size of a file in the gateway cache. Units: Bytes
`CacheFree`	This metric reports the number of available bytes in the gateway cache. Units: Bytes
`CacheHitPercent`	Percent of application read operations from the gateway that are served from cache. The sample is taken at the end of the reporting period. When there are no application read operation s from the gateway, this metric reports 100 percent. Units: Percent

Metric	Description
CachePercentDirty	The overall percentage of the gateway cache that has not been persisted to AWS. The sample is taken at the end of the reporting period. Units: Percent
CachePercentUsed	The overall percent of the gateway cache storage that is used. The sample is taken at the end of the reporting period. Units: Percent
CacheUsed	This metric reports the number of used bytes in the gateway cache. Units: Bytes
CloudBytesDownloaded	The total number of bytes that the gateway downloaded from AWS during the reporting period. Use this metric with the Sum statistic to measure throughput and with the Samples statistic to measure IOPS. Units: Bytes
CloudBytesUploaded	The total number of bytes that the gateway uploaded to AWS during the reporting period. Use this metric with the Sum statistic to measure throughput and with the Samples statistic to measure input/output operations per second (IOPS). Units: Bytes

Metric	Description
FilesFailingUpload	This metric tracks the number of files which are failing to upload to AWS. These files will generate health notifications which contain more information on the issue. Use this metric with the Sum statistic to show the number of files which are currently failing to upload to AWS. Units: Count
FileShares	This metric reports the number of file shares on the gateway. Units: Count
FileSystem-ERROR	This metric provides the number of file system associations on this gateways which are in the ERROR state. If this metric reports any file system associati ons are in the ERROR state, then it is likely there is a problem with the gateway which is may cause disruption to your workflow. It is recommended to create an alarm for when this metric reports a non-zero value. Units: Count
HealthNotifications	This metric reports the number of health notifications that were generated by this gateway in the reporting period. Units: Count

Metric	Description
IndexEvictions	This metric reports the number of files whose metadata was evicted from the cached index of file metadata to make room for new entries. The gateway maintains this metadata index, which is populated from the AWS Cloud on demand. Units: Count
IndexFetches	This metric reports the number of files for which metadata was fetched. The gateway maintains a cached index of file metadata, which is populated from the AWS Cloud on demand. Units: Count
IoWaitPercent	This metric reports the percentage of time that the CPU is waiting for a response from the local disk. Units: Percent
MemTotalBytes	This metric reports the total amount of memory on the gateway. Units: Bytes
MemUsedBytes	This metric reports the amount of used memory on the gateway. Units: Bytes

Metric	Description
RootDiskFreeBytes	This metric reports the number of available bytes on the root disk of the gateway. If this metric reports less than 20 GB are free, you should increase the size of the root disk. To increase the root disk size, you can increase the size of existing root disk on the VM. When the VM is rebooted, gateway recognizes the increased size on the root disk. Units: Bytes
SmbV2Sessions	This metric reports the number of SMBv2 sessions that are active on the gateway. This metric is emitted once for each file system associated with the gateway. Use the SUM stat to calculate the total number of active SMBv2 sessions across all file systems. Units: Count
SmbV3Sessions	This metric reports the number of SMBv3 sessions that are active on the gateway. This metric is emitted once for each file system associated with the gateway. Use the SUM stat to calculate the total number of active SMBv3 sessions across all file systems. Units: Count
TotalCacheSize	This metric reports the total size of the cache. Units: Bytes

Metric	Description
`UserCpuPercent`	This metric reports the percentage of time that is spent on gateway processing. Units: Percent

Understanding file system metrics

You can find information following about the Storage Gateway metrics that cover file systems. Each file system has a set of metrics associated with it. Some file system-specific metrics have the same name as certain gateway-specific metrics. These metrics represent the same kinds of measurements, but are scoped to the file system instead.

Always specify whether you want to work with either a gateway or a file system metric before working with a metric. Specifically, when working with file system metrics, you must specify the `File system ID` that identifies the file system for which you are interested in viewing metrics. For more information, see Using Amazon CloudWatch metrics.

> **Note**
>
> Some metrics return data points only when new data has been generated during the most recent monitoring period.

The following table describes the Storage Gateway metrics that you can use to get information about your file shares.

Metric	Description
`CacheHitPercent`	Percent of application read operations from the file shares that are served from cache. The sample is taken at the end of the reporting period. When there are no application read operation s from the file share, this metric reports 100 percent.

Metric	Description
	Units: Percent
CachePercentDirty	The file share's contribution to the overall percentage of the gateway's cache that has not been persisted to AWS. The sample is taken at the end of the reporting period. Use this metric with the Sum statistic. Ideally, this metric should remain low. **Note** Use the CachePercentDirty metric of the gateway to view the overall percentage of the gateway's cache that has not been persisted to AWS. Units: Percent
CachePercentUsed	The percent of the data cache used across the entire gateway. The sample is taken at the end of the reporting period. This file share-specific metric reports the same value as the corresponding gateway-specific metric. Units: Percent
CloudBytesUploaded	The total number of bytes that the gateway uploaded to AWS during the reporting period. Use this metric with the Sum statistic to measure throughput and with the Samples statistic to measure IOPS. Units: Bytes

Metric	Description
CloudBytesDownloaded	The total number of bytes that the gateway downloaded from AWS during the reporting period. Use this metric with the Sum statistic to measure throughput and with the Samples statistic to measure input/output operations per second (IOPS). Units: Bytes
FilesFailingUpload	This metric tracks the number of files which are failing to upload to AWS. These files will generate health notifications which contain more information on the issue. Use this metric with the Sum statistic to show the number of files which are currently failing to upload to AWS. Units: Count
ReadBytes	The total number of bytes read from your on-premises applications in the reporting period for a file share. Use this metric with the Sum statistic to measure throughput and with the Samples statistic to measure IOPS. Units: Bytes

Metric	Description
`WriteBytes`	The total number of bytes written to your on-premises applications in the reporting period. Use this metric with the `Sum` statistic to measure throughput and with the `Samples` statistic to measure IOPS. Units: Bytes

Understanding FSx File Gateway audit logs

Amazon FSx File Gateway (FSx File Gateway) audit logs provide you with details about user access to files and folders within a file system association. You can use audit logs to monitor user activities and take action if inappropriate activity patterns are identified. The logs are formatted similar to Windows Server security log events, to support compatibility with existing log processing tools for Windows security events.

Operations

The following table describes the FSx File Gateway audit log file access operations.

Operation name	Definition
Read Data	Read the contents of a file.
Write Data	Change the contents of a file.
Create	Create a new file or folder.
Rename	Rename an existing file or folder.
Delete	Delete a file or folder.
Write Attributes	Update file or folder metadata (ACLs, owner, group, permissions).

Attributes

The following table describes FSx File Gateway audit log file access attributes.

Attribute	Definition
`securityDescriptor`	Shows the discretionary access control list (DACL) set on an object, in SDDL format.
`sourceAddress`	The IP address of file share client machine.
`SubjectDomainName`	The Active Directory (AD) domain that the client's account belongs to.
`SubjectUserName`	The Active Directory user name of the client.
`source`	The ID of the Storage Gateway `FileSystemAssociation` that is being audited.
`mtime`	This time that the object's content was modified, set by the client.
`version`	The version of the audit log format.
`ObjectType`	Defines whether the object is a file or folder.
`locationDnsName`	The FSx File Gateway system DNS name.
`objectName`	The full path to the object.
`ctime`	The time that the object's content or metadata was modified, set by the client.
`shareName`	The name of the share that is being accessed.
`operation`	The name of the object access operation.
`newObjectName`	The full path to the new object after it has been renamed.
`gateway`	The Storage Gateway ID.

Attribute	Definition
status	The status of the operation. Only success is logged (failures are logged with the exception of failures arising from permissions denied).
fileSizeInBytes	The size of the file in bytes, set by the client at file creation time.

Attributes logged per operation

The following table describes the FSx File Gateway audit log attributes logged in each file access operation.

	Read data	Write data	Create folder	Create file	Rename file/ folder	Delete file/ folder	Write attributes (change ACL)	Write attributes (chown)	Write attributes (chmod)	Write attributes (chgrp)
securiescrip							X			
sourceress	X	X	X	X	X	X	X	X	X	X
SubjecmainNa	X	X	X	X	X	X	X	X	X	X
SubjecerName	X	X	X	X	X	X	X	X	X	X
source	X	X	X	X	X	X	X	X	X	X
mtime			X	X						
versic	X	X	X	X	X	X	X	X	X	X

	Read data	Write data	Create folder	Create file	Rename file/ folder	Delete file/ folder	Write attributes (change ACL)	Write attributes (chown)	Write attributes (chmod)	Write attributes (chgrp)
object	X	X	X	X	X	X	X	X	X	X
locationsName	X	X	X	X	X	X	X	X	X	X
object	X	X	X	X	X	X	X	X	X	X
ctime			X	X						
shareN	X	X	X	X	X	X	X	X	X	X
operat	X	X	X	X	X	X	X	X	X	X
newObjName					X					
gatewa	X	X	X	X	X	X	X	X	X	X
status	X	X	X	X	X	X	X	X	X	X
fileSinBytes				X						

Maintaining your gateway

Maintaining your Amazon FSx File Gateway involves doing general maintenance to optimize your gateway's performance. These tasks are common to all gateway types.

This section contains the following topics, which describe concepts and procedures related to maintaining your Amazon FSx File Gateway:

Topics

- Managing gateway updates - Learn how to turn maintenance updates on or off, and modify the maintenance window schedule for your File Gateway.
- Performing maintenance tasks using the local console - Learn how to perform maintenance tasks using the gateway local console.
- Shutting down your gateway VM - Learn about what to do if you need to shutdown or reboot your gateway virtual machine for maintenance, such as when applying a patch to your hypervisor.
- Replacing your existing FSx File Gateway with a new instance - Learn how to replace your FSx File Gateway with a new instance when you want to improve performance or to respond to a notification to migrate the gateway.
- Deleting your gateway and removing associated resources - Learn how to delete your gateway using the AWS Storage Gateway console and clean up associated resources to avoid being charged for their continued use.

Managing gateway updates

Storage Gateway consists of a managed cloud services component and a gateway appliance component that you deploy either on-premises, or on an Amazon EC2 instance in the AWS cloud. Both components receive regular updates. The topics in this section describe the cadence of these updates, how they are applied, and how to configure update-related settings on the gateways in your deployment.

> **Important**
>
> You should treat the Storage Gateway appliance as a managed virtual machine, and should not attempt to access or modify its installation in any way. Attempting to install or

update any software packages using methods other than the normal AWS gateway update mechanism (for example, SSM or hypervisor tools) may cause the gateway to malfunction.

Update frequency and expected behavior

AWS updates the cloud services component as needed without causing disruption to deployed gateways. Your deployed gateway appliances receive the following types of updates:

- Maintenance - Regular updates that can include operating system and software upgrades, fixes to address stability, performance, and security, and access to new features.
- Urgent - Critical updates that include required fixes for issues that immediately impact the security, performance, or durability of your gateway. Urgent updates can be released at any time, outside the normal cadence of the monthly maintenance and feature updates.

All updates are cumulative, and upgrade gateways to the current version when applied. For information about the specific changes included in each update, see .

All gateway appliance updates may cause a brief disruption of service. The gateway's VM host doesn't need to reboot during updates, but the gateway will be unavailable for a short period while the gateway appliance updates and restarts.

When you deploy and activate your gateway, a default maintenance window schedule is set. You can modify the maintenance window schedule at any time. You can also turn off maintenance updates, but we recommend leaving them turned on.

Note

Urgent updates will be applied according to the maintenance window schedule, even if regular maintenance updates are turned off.

Before any update is applied to your gateway, AWS notifies you with a message on the Storage Gateway console and your AWS Health Dashboard. For more information, see AWS Health Dashboard. To modify the email address where software update notifications are sent, see Update the alternate contacts for your AWS account in the *AWS Account Management Reference Guide*.

When updates are available, the gateway **Details** tab displays a maintenance message. You can also see the date and time that the last successful update was applied on the **Details** tab.

Turn maintenance updates on or off

When maintenance updates are turned on, your gateway automatically applies these updates according to the configured maintenance window schedule. For more information, see Modify the gateway maintenance window schedule.

If maintenance updates are turned off, the gateway will not apply these updates automatically, but you can always apply them manually using the Storage Gateway console, API, or CLI. Urgent updates will sometimes be applied during your configured maintenance window, regardless of this setting.

> **Note**
>
> The following procedure describes how to turn gateway updates on or off using the Storage Gateway console. To change this setting programmatically using the API, see UpdateMaintenanceStartTime in the *Storage Gateway API Reference*.

To turn maintenance updates on or off using the Storage Gateway console:

1. Open the Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. On the navigation pane, choose **Gateways**, and then choose the gateway for which you want to configure maintenance updates.
3. Choose **Actions**, and then choose **Edit maintenance settings**.
4. For **Maintenance updates**, select **On** or **Off**.
5. Choose **Save changes** when finished.

You can verify the updated setting on the **Details** tab for the selected gateway in the Storage Gateway console.

Modify the gateway maintenance window schedule

If maintenance updates are turned on, your gateway automatically applies these updates according the maintenance window schedule. Urgent updates will sometimes be applied during your configured maintenance window, regardless of the maintenance updates setting.

> **Note**
>
> The following procedure describes how to modify the maintenance window schedule using the Storage Gateway console. To change this setting programmatically using the API, see UpdateMaintenanceStartTime in the *Storage Gateway API Reference*.

To modify the maintenance window schedule using the Storage Gateway console:

1. Open the Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. On the navigation pane, choose **Gateways**, and then choose the gateway for which you want to configure maintenance updates.
3. Choose **Actions**, and then choose **Edit maintenance settings**.
4. Under **Maintenance window start time**, do the following:

   a. For **Schedule**, choose **Weekly** or **Monthly** to set the maintenance window cadence.

   b. If you choose **Weekly**, modify the values for **Day of the week** and **Time** to set the specific point during each week when the maintenance window will begin.

      If you choose **Monthly**, modify the values for **Day of the month** and **Time** to set the specific point during each month when the maintenance window will begin.

      > **Note**
      >
      > The maximum value that can be set for day of the month is 28. It is not possible to set the maintenance schedule to start on days 29 through 31.
      > If you receive an error while configuring this setting, it might mean that your gateway software is out of date. Considering updating your gateway manually first, and then attempt to configure the maintenance window schedule again.

5. Choose **Save changes** when finished.

You can verify the updated settings on the **Details** tab for the selected gateway in the Storage Gateway console.

Apply an update manually

If a software update is available for your gateway, you can apply it manually by following the procedure below. This manual update process ignores the maintenance window schedule and applies the update immediately, even if maintenance updates are turned off.

> **Note**
>
> The following procedure describes how to manually apply an update using the Storage Gateway console. To perform this action programmatically using the API, see UpdateGatewaySoftwareNow in the *Storage Gateway API Reference*.

To apply a gateway software update manually using the Storage Gateway console:

1. Open the Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. On the navigation pane, choose **Gateways**, and then choose the gateway you want to update.

   If an update is available, the console displays a blue notification banner on the gateway **Details** tab, which includes an option to apply the update.
3. Choose **Apply update now** to immediately update the gateway.

   > **Note**
   >
   > This operation causes a temporary disruption to gateway functionality while the update installs. During this time, the gateway status appears **OFFLINE** in the Storage Gateway console. After the update finishes installing, the gateway resumes normal operation and its status changes to **RUNNING**.

You can verify that the gateway software was updated to the latest version by checking the **Details** tab for the selected gateway in the Storage Gateway console.

Performing maintenance tasks using the local console

This section contains the following topics, which provide information about how to perform maintenance tasks using the gateway appliance local console. You can perform these tasks by accessing the local console through the on-premises virtual machine or Amazon EC2 instance that hosts your gateway appliance. Most of the tasks are common across the different host platforms, but there are also some differences.

Topics

- Accessing the gateway local console - Learn how to log into the local console for an on-premises gateway hosted on a Linux Kernel-based Virtual Machine (KVM), VMware ESXi, or Microsoft Hyper-V Manager platform.
- Performing tasks on the virtual machine local console - Learn how to use the local console to perform basic setup and advanced configuration tasks for an on-premises gateway, such as configuring an HTTP proxy, viewing system resource status, or running terminal commands.
- Performing tasks on the Amazon EC2 gateway local console - Learn how to log into the local console to perform basic setup and advanced configuration tasks for an Amazon EC2 gateway, such as configuring an HTTP proxy, viewing system resource status, or running terminal commands.

Accessing the gateway local console

How you access your VM's local console depends on the type of the Hypervisor you deployed your gateway VM on. In this section, you can find information on how to access the VM local console using Linux Kernel-based Virtual Machine (KVM), VMware ESXi, and Microsoft Hyper-V Manager.

Topics

- Accessing the Gateway Local Console with Linux KVM
- Accessing the Gateway Local Console with VMware ESXi
- Access the Gateway Local Console with Microsoft Hyper-V

Accessing the Gateway Local Console with Linux KVM

There are different ways to configure virtual machines running on KVM, depending on the Linux distribution being used. Instructions for accessing KVM configuration options from the command line follow. Instructions might differ depending on your KVM implementation.

To access your gateway's local console with KVM

1. Use the following command to list the VMs that are currently available in KVM.

   ```
   # virsh list
   ```

   The command returns a list of VMs with **Id**, **Name**, and **State** information for each. Note the `Id` of the VM for which you want to launch the gateway local console.

2. Use the following command to access the local console.

   ```
   # virsh console Id
   ```

   Replace *Id* with the **Id** of the VM you noted in the previous step.

   The AWS Appliance gateway local console prompts you to login to change your network configuration and other settings.

3. Enter your username and password to log into the gateway local console. For more information, see Logging in to the File Gateway local console .

   After you log in, the **AWS Appliance Activation - Configuration** menu appears. You can select from the menu options to perform gateway configuration tasks. For more information, see Performing tasks on the virtual machine local console .

Accessing the Gateway Local Console with VMware ESXi

To access your gateway's local console with VMware ESXi

1. In the VMware vSphere client, select your gateway VM.
2. Make sure that the gateway VM is turned on.

> **Note**
>
> If your gateway VM is turned on, a green arrow icon appears with the VM icon in the VM browser panel on the left side of the application window. If your gateway VM is not turned on, you can turn it on by choosing the green **Power On** icon on the **Toolbar** at the top of the application window.

3. Choose the **Console** tab in the main information panel on the right side of the application window.

   After a few moments, the AWS Appliance gateway local console prompts you to login to change your network configuration and other settings.

   > **Note**
   >
   > To release the cursor from the console window, press **Ctrl+Alt**.

4. Enter your username and password to log into the gateway local console. For more information, see Logging in to the File Gateway local console .

   After you log in, the **AWS Appliance Activation - Configuration** menu appears. You can select from the menu options to perform gateway configuration tasks. For more information, see Performing tasks on the virtual machine local console .

Access the Gateway Local Console with Microsoft Hyper-V

To access your gateway's local console (Microsoft Hyper-V)

1. Select your gateway appliance VM from the **Virtual Machines** panel on the left side of the Microsoft Hyper-V Manager application window.
2. Make sure that the gateway is turned on.

   > **Note**
   >
   > If your gateway VM is turned on, `Running` is displayed in the **State** column for the VM in the **Virtual Machines** panel on the left side of the application window. If your

gateway VM is not turned on, you can turn it on by choosing **Start** in the **Actions** panel on the right side of the application window.

3. Choose **Connect** from the **Actions** panel.

   The **Virtual Machine Connection** window appears. If an authentication window appears, type the sign-in credentials provided to you by the hypervisor administrator.

   After a few moments, the AWS Appliance gateway local console prompts you to login to change your network configuration and other settings.

4. Enter your username and password to log into the gateway local console. For more information, see Logging in to the File Gateway local console .

   After you log in, the **AWS Appliance Activation - Configuration** menu appears. You can select from the menu options to perform gateway configuration tasks. For more information, see Performing tasks on the virtual machine local console .

Performing tasks on the virtual machine local console

For a File Gateway deployed on-premises, you can perform the following maintenance tasks using the VM host's local console. These tasks are common to VMware, Microsoft Hyper-V, and Linux Kernel-based Virtual Machine (KVM) hypervisors.

Topics

- Logging in to the File Gateway local console - Learn how to login to the local console where you can configure gateway network settings and change the default password.
- Configuring an HTTP proxy - Learn how to configure Storage Gateway to route all AWS endpoint traffic through a proxy server.
- Configuring your gateway network settings - Learn how to configure your gateway to use DHCP or a static IP address.
- Testing your gateway's network connectivity - Learn how to use the gateway local console to test network connectivity.
- Viewing your gateway system resource status - Learn how to check your gateway's virtual CPU cores, root volume size, and RAM.

- Configuring a Network Time Protocol (NTP) server for your gateway - Learn how to view and edit Network Time Protocol (NTP) server configurations and synchronize the time on your gateway with your hypervisor host.
- Running Storage Gateway commands on the local console - Learn how to run local console commands to perform tasks such as saving routing tables, connecting to Support, and more.

Logging in to the File Gateway local console

When the VM is ready for you to log in, the login screen is displayed. If this is your first time logging in to the VM local console, you use the temporary sign-in credentials to log in. These temporary credentials give you access to menus where you can configure gateway network settings and change the password from the local console. The initial user name is `admin` and the temporary password is `password`. You must change the password on first log in.

To change the temporary password

1. On the **AWS Appliance Activation - Configuration** main menu, enter the corresponding numeral for **Gateway Console**.
2. Run the `passwd` command. For information about how to run the command, see Running Storage Gateway commands on the local console.

Setting the local console password from the Storage Gateway console

You can also manage the local console's password from the Storage Gateway web-based console. Any successful password updates made with the web-based console will override the password used by the gateway VM's local console, including the temporary password if you have never logged in locally. If the gateway is not currently reachable over the network, the password update process will fail.

To set the local console password on the Storage Gateway console

1. Open the Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. On the navigation pane, choose **Gateways**, and then select the gateway for which you want to set a new password.
3. For **Actions**, choose **Set Local Console Password**.

4. In the **Set Local Console Password** dialog box, enter a new password, confirm the password, and then choose **Save**.

   Your new password replaces the current password. The Storage Gateway service doesn't save, store, or log the password but instead safely transmits it over an encrypted channel to the VM, where it is securely stored.

   > **Note**
   >
   > The password can consist of any character on the keyboard and can be 1–512 characters long.

Configuring an HTTP proxy

File Gateways support configuration of an HTTP proxy.

> **Note**
>
> The only proxy configuration that File Gateways support is HTTP.

If your gateway must use a proxy server to communicate to the internet, then you need to configure the HTTP proxy settings for your gateway. You do this by specifying an IP address and port number for the host running your proxy. After you do so, Storage Gateway routes all AWS endpoint traffic through your proxy server. Communications between the gateway and endpoints is encrypted, even when using the HTTP proxy. For information about network requirements for your gateway, see Network and firewall requirements.

To configure an HTTP proxy for a File Gateway

1. Log in to your gateway's local console:

   - For more information on logging in to the VMware ESXi local console, see Accessing the Gateway Local Console with VMware ESXi.
   - For more information on logging in to the Microsoft Hyper-V local console, see Access the Gateway Local Console with Microsoft Hyper-V.
   - For more information on logging in to the local console for the Linux Kernel-Based Virtual Machine (KVM), see Accessing the Gateway Local Console with Linux KVM.

2. From the **AWS Appliance Activation - Configuration** main menu, enter the corresponding numeral to select **Configure HTTP Proxy**.
3. From the **AWS Appliance Activation HTTP Proxy Configuration** menu, enter the corresponding numeral for the task you want to perform:
   - **Configure HTTP proxy** - You will need to supply a host name and port to complete configuration.
   - **View current HTTP proxy configuration** - If an HTTP proxy is not configured, the message `HTTP Proxy not configured` is displayed. If an HTTP proxy is configured, the host name and port of the proxy are displayed.
   - **Remove an HTTP proxy configuration** - The message `HTTP Proxy Configuration Removed` is displayed.
4. Restart your VM to apply your HTTP configuration settings.

Configuring your gateway network settings

The default network configuration for the gateway is Dynamic Host Configuration Protocol (DHCP). With DHCP, your gateway is automatically assigned an IP address. In some cases, you might need to manually assign your gateway's IP as a static IP address, as described following.

To configure your gateway to use static IP addresses

1. Log in to your gateway's local console:
   - For more information on logging in to the VMware ESXi local console, see Accessing the Gateway Local Console with VMware ESXi.
   - For more information on logging in to the Microsoft Hyper-V local console, see Access the Gateway Local Console with Microsoft Hyper-V.
   - For more information on logging in to the KVM local console, see Accessing the Gateway Local Console with Linux KVM.
2. From the **AWS Appliance Activation - Configuration** main menu, enter the corresponding numeral to select **Network Configuration**.
3. From the **Network Configuration** menu, perform one of the following tasks:

To Perform This Task	Do This
Get information about your network adapter	Enter the corresponding numeral to select **Describe Adapter**. A list of adapter names appears, and you are prompted to enter an adapter name—for example, **`eth0`**. If the adapter you specify is in use, the following information about the adapter is displayed: • Media access control (MAC) address • IP address • Netmask • Gateway IP address • DHCP enabled status You use the adapter names listed here when you configure a static IP address or when you set your gateway's default adapter.
Configure DHCP routing	Enter the corresponding numeral to select **Configure DHCP**. You are prompted to configure the network interface to use DHCP.

To Perform This Task	Do This
Configure a static IP address for your gateway	Enter the corresponding numeral to select **Configure Static IP**. You are prompted to enter the following information to configure a static IP: • Network adapter name • IP address • Netmask • Default gateway address • Primary Domain Name Service (DNS) address • Secondary DNS address **Important** If your gateway has already been activated, you must shut it down and restart it from the Storage Gateway console for the settings to take effect. For more information, see Shutting down your gateway VM. If your gateway uses more than one network interface, you must set all active interfaces to use DHCP or static IP addresses.

To Perform This Task	Do This
	For example, suppose that your gateway VM uses two interfaces configured as DHCP. If you later set one interface to a static IP, the other interface is deactivated. To activate the interface in this case, you must set it to a static IP. If both interfaces are initially set to use static IP addresses and you then set the gateway to use DHCP, both interfaces use DHCP.
Configure a hostname for your gateway	Enter the corresponding numeral to select **Configure Hostname**. You are prompted to choose whether the gateway will use a static hostname that you specify, or aquire one automatically through DCHP or rDNS. If you select **Static**, you are prompted to provide a static hostname, such as `testgateway.example.com` . Enter `y` to apply the configuration. **Note** If you configure a static hostname for your gateway, ensure that the provided hostname is in the domain that gateway is joined to. You must also create an A record in your DNS system that points the gateway's IP address to its static hostname.

To Perform This Task	Do This
View your gateway's hostname configura tion	Enter the corresponding numeral to select **View Hostname Configuration**. Your gateway's hostname, aquisition mode, domain, and Active Directory realm are displayed.
Reset all your gateway's network configuration to DHCP	Enter the corresponding numeral to select **Reset all to DHCP**. All network interfaces are set to use DHCP. **Important** If your gateway has already been activated, you must shut down and restart your gateway from the Storage Gateway console for the settings to take effect. For more information, see Shutting down your gateway VM.
Set your gateway's default route adapter	Enter the corresponding numeral to select **Set Default Adapter**. The available adapters for your gateway are shown, and you are prompted to choose one of the adapters—for example, **`eth0`**.

To Perform This Task	Do This
Edit your gateway's DNS configuration	Enter the corresponding numeral to select **Edit DNS Configuration**. The available adapters of the primary and secondary DNS servers are displayed. You are prompted to provide the new IP address.
View your gateway's DNS configuration	Enter the corresponding numeral to select **View DNS Configuration**. The available adapters of the primary and secondary DNS servers are displayed. **Note** For some versions of the VMware hypervisor, you can edit the adapter configuration in this menu.
View routing tables	Enter the corresponding numeral to select **View Routes**. The default route of your gateway is displayed.

Testing your gateway's network connectivity

You can use your gateway's local console to test your network connectivity. This test can be useful when you are troubleshooting network issues with your gateway.

To test your gateway's network connectivity

1. Log in to your gateway's local console:

- For more information on logging in to the VMware ESXi local console, see Accessing the Gateway Local Console with VMware ESXi.
- For more information on logging in to the Microsoft Hyper-V local console, see Access the Gateway Local Console with Microsoft Hyper-V.
- For more information on logging in to the KVM local console, see Accessing the Gateway Local Console with Linux KVM.

2. From the **AWS Appliance Activation - Configuration** main menu, enter the corresponding numeral to select **Test Network Connectivity**.

   If your gateway has already been activated, the connectivity test begins immediately. For gateways that have not yet been activated, you must specify the endpoint type and AWS Region as described in the following steps.

3. If your gateway is not yet activated, enter the corresponding numeral to select the endpoint type for your gateway.
4. If you selected the public endpoint type, enter the corresponding numeral to select the AWS Region that you want to test. For supported AWS Regions and a list of AWS service endpoints you can use with Storage Gateway, see AWS Storage Gateway endpoints and quotas in the *AWS General Reference*.

As the test progresses, each endpoint displays either **[PASSED]** or **[FAILED]**, indicating the status of the connection as follows:

Message	Description
[PASSED]	Storage Gateway has network connectivity.
[FAILED]	Storage Gateway does not have network connectivity.

Viewing your gateway system resource status

When your gateway starts, it checks its virtual CPU cores, root volume size, and RAM. It then determines whether these system resources are sufficient for your gateway to function properly. You can view the results of this check on the gateway's local console.

To view the status of a system resource check

1. Log in to your gateway's local console:

   - For more information on logging in to the VMware ESXi console, see Accessing the Gateway Local Console with VMware ESXi.
   - For more information on logging in to the Microsoft Hyper-V local console, see Access the Gateway Local Console with Microsoft Hyper-V.
   - For more information on logging in to the KVM local console, see Accessing the Gateway Local Console with Linux KVM.

2. From the **AWS Appliance Activation - Configuration** main menu, enter the corresponding numeral to select **View System Resource Check**.

   Each resource displays **[OK]**, **[WARNING]**, or **[FAIL]**, indicating the status of the resource as follows:

Message	Description
[OK]	The resource has passed the system resource check.
[WARNING]	The resource doesn't meet the recommended requirements, but your gateway can continue to function. Storage Gateway displays a message that describes the results of the resource check.
[FAIL]	The resource doesn't meet the minimum requirements. Your gateway might not function properly. Storage Gateway displays a message that describes the results of the resource check.

   The console also displays the number of errors and warnings next to the resource check menu option.

Configuring a Network Time Protocol (NTP) server for your gateway

You can view and edit Network Time Protocol (NTP) server configurations and synchronize the VM time on your gateway with your hypervisor host.

To manage system time

1. Log in to your gateway's local console:
   - For more information on logging in to the VMware ESXi local console, see Accessing the Gateway Local Console with VMware ESXi.
   - For more information on logging in to the Microsoft Hyper-V local console, see Access the Gateway Local Console with Microsoft Hyper-V.
   - For more information on logging in to the KVM local console, see Accessing the Gateway Local Console with Linux KVM.
2. From the **AWS Appliance Activation - Configuration** main menu, enter the corresponding numeral to select **System Time Management**.
3. From the **System Time Management** menu, enter the corresponding numeral to perform one of the following tasks.

To Perform This Task	Do This
View and synchronize your VM time with NTP server time.	Enter the corresponding numeral to select **View and Synchronize System Time**. The current time of your VM is displayed . Your File Gateway determines the time difference from your gateway VM, and your NTP server time prompts you to synchronize the VM time with NTP time. After your gateway is deployed and running, in some scenarios the gateway VM's time can drift. For example, suppose that there is a prolonged network outage and your hypervisor host and gateway don't get time updates. In this case, the gateway VM's time

To Perform This Task	Do This
	is different from the true time. When there is a time drift, a discrepancy occurs between the stated times when operations such as snapshots occur and the actual times that the operations occur. For a gateway deployed on VMware ESXi, setting the hypervisor host time and synchronizing the VM time to the host is sufficient to avoid time drift. For more inf ormation, see Synchronize VM time with VMware host time. For a gateway deployed on Microsoft Hyper-V, you should periodically check your VM's time. For more information, see Synchronize VM time with Hyper-V or Linux KVM host time. For a gateway deployed on KVM, you can check and synchronize the VM time using `virsh` command line interface for KVM.
Edit your NTP server configuration	Enter the corresponding numeral to select **Edit NTP Configuration**. You are prompted to provide a preferred and a secondary NTP server.
View your NTP server configuration	Enter the corresponding numeral to select **View NTP Configuration**. Your NTP server configuration is displayed.

Running Storage Gateway commands on the local console

The VM local console in Storage Gateway helps provide a secure environment for configuring and diagnosing issues with your gateway. Using the local console commands, you can perform maintenance tasks such as saving routing tables, connecting to Support, and so on.

To run a configuration or diagnostic command

1. Log in to your gateway's local console:
   - For more information on logging in to the VMware ESXi local console, see Accessing the Gateway Local Console with VMware ESXi.
   - For more information on logging in to the Microsoft Hyper-V local console, see Access the Gateway Local Console with Microsoft Hyper-V.
   - For more information on logging in to the KVM local console, see Accessing the Gateway Local Console with Linux KVM.
2. From the **AWS Appliance Activation - Configuration** main menu, enter the corresponding numeral to select **Gateway Console**.
3. From the gateway console command prompt, enter **h**.

   The console displays the **AVAILABLE COMMANDS** menu, which lists the available commands:

Command	Function
dig	Collect output from dig for DNS troublesh ooting.
exit	Return to Configuration menu.
h	Display available command list.
ifconfig	View or configure network interfaces. **Note** We recommend configuring network or IP settings using the Storage Gateway console or the dedicated

Command	Function
	local console menu option. For instructions, see Configuring your gateway network settings.
ip	Show / manipulate routing, devices, and tunnels. **Note** We recommend configuring network or IP settings using the Storage Gateway console or the dedicated local console menu option. For instructions, see Configuring your gateway network settings.
iptables	Administration tool for IPv4 packet filtering and NAT.
ncport	Test connectivity to a specific TCP port on a network.
nping	Collect output from nping for network troubleshooting.
open-support-channel	Connect to AWS Support. For instructions on how to turn on AWS support access, see You want AWS Support to help troubleshoot your EC2 gateway.
passwd	Update authentication tokens.
save-iptables	Persist IP tables.
save-routing-table	Save newly added routing table entry.

Command	Function
tcptraceroute	Collect traceroute output on TCP traffic to a destination.
sslcheck	Returns output with certificate issuer **Note** Storage Gateway uses certificate issuer verification and does not support ssl inspection. If this command returns an issuer other than aws-appliance@amazon.com, then it is likely that an application performing an ssl inspection. In that case, we recommend bypassing ssl inspection for the Storage Gateway appliance.

4. From the gateway console command prompt, enter the corresponding command for the function you want to use, and follow the instructions.

To learn about a command, enter **man** + *command name* at the command prompt.

Performing tasks on the Amazon EC2 gateway local console

Some maintenance tasks require that you log in to the local console when running a gateway deployed on an Amazon EC2 instance. This section describes how to log in to the local console and perform maintenance tasks.

Topics

- Logging in to your Amazon EC2 gateway local console - Learn how to connect and log in to the gateway local console your Amazon EC2 instance by using a Secure Shell (SSH) client.
- Routing your gateway deployed on Amazon EC2 through an HTTP proxy - Learn how to configure a Socket Secure version 5 (SOCKS5) proxy between AWS and a gateway deployed on an Amazon EC2 instance.

- Testing your gateway's network connectivity - Learn how to use the gateway local console to test network connectivity between your gateway and various network resources.
- Viewing your gateway system resource status - Learn how to use the gateway local console to checks your gateway's virtual CPU cores, root volume size, and RAM.
- Running Storage Gateway commands on the local console for an Amazon EC2 gateway - Learn how to run local console commands to perform tasks such as saving routing tables, connecting to Support, and more.
- Configuring your Amazon EC2 gateway network settings - Learn how to use the local console to view and configure network settings such as DNS and hostname for a gateway on an Amazon EC2 instance.

Logging in to your Amazon EC2 gateway local console

You log in to the gateway local console on an Amazon EC2 instance by using a Secure Shell (SSH) client. For detailed information, see Connect to your instance in the *Amazon EC2 User Guide*. To connect this way, you need the SSH key pair that you specified when you launched your instance. For information about Amazon EC2 key pairs, see Amazon EC2 key pairs in the *Amazon EC2 User Guide.*

To log in to the gateway local console

1. Connect to the Amazon EC2 instance using SSH and log in as the *admin* user.
2. After you log in, you see the **AWS Appliance Activation - Configuration** main menu, from which you can perform various tasks.

To Learn About This Task	See This Topic
Configure an HTTP proxy for your gateway	Routing your gateway deployed on Amazon EC2 through an HTTP proxy
Configure network settings for your gateway	Configuring your Amazon EC2 gateway network settings
Test network connectivity	Testing your gateway's network connectivity
View a system resource check	Viewing your gateway system resource status.

To Learn About This Task	See This Topic
Run Storage Gateway console commands	Running Storage Gateway commands on the local console for an Amazon EC2 gateway

To shut down the gateway, enter **0**.

To exit the configuration session, enter **X**.

Routing your gateway deployed on Amazon EC2 through an HTTP proxy

Storage Gateway supports the configuration of a Socket Secure version 5 (SOCKS5) proxy between your gateway deployed on Amazon EC2 and AWS.

If your gateway must use a proxy server to communicate to the internet, then you need to configure the HTTP proxy settings for your gateway. You do this by specifying an IP address and port number for the host running your proxy. After you do so, Storage Gateway routes all AWS endpoint traffic through your proxy server. Communications between the gateway and endpoints is encrypted, even when using the HTTP proxy.

To route your gateway internet traffic through a local proxy server

1. Log in to your gateway's local console. For instructions, see Logging in to your Amazon EC2 gateway local console.
2. From the **AWS Appliance Activation - Configuration** main menu, enter the corresponding numeral to select **Configure HTTP Proxy**.
3. From the **AWS Appliance Activation HTTP Proxy Configuration** menu, enter the corresponding numeral for the task you want to perform:

   - **Configure HTTP proxy** - You will need to supply a host name and port to complete configuration.
   - **View current HTTP proxy configuration** - If an HTTP proxy is not configured, the message `HTTP Proxy not configured` is displayed. If an HTTP proxy is configured, the host name and port of the proxy are displayed.
   - **Remove an HTTP proxy configuration** - The message `HTTP Proxy Configuration Removed` is displayed.

Testing your gateway's network connectivity

You can use your gateway's local console to test your network connectivity. This test can be useful when you are troubleshooting network issues with your gateway.

To test your gateway's connectivity

1. Log in to your gateway's local console. For instructions, see Logging in to your Amazon EC2 gateway local console.
2. From the **AWS Appliance Activation - Configuration** main menu, enter the corresponding numeral to select **Test Network Connectivity**.

   If your gateway has already been activated, the connectivity test begins immediately. For gateways that have not yet been activated, you must specify the endpoint type and AWS Region as described in the following steps.
3. If your gateway is not yet activated, enter the corresponding numeral to select the endpoint type for your gateway.
4. If you selected the public endpoint type, enter the corresponding numeral to select the AWS Region that you want to test. For supported AWS Regions and a list of AWS service endpoints you can use with Storage Gateway, see AWS Storage Gateway endpoints and quotas in the *AWS General Reference*.

As the test progresses, each endpoint displays either **[PASSED]** or **[FAILED]**, indicating the status of the connection as follows:

Message	Description
[PASSED]	Storage Gateway has network connectivity.
[FAILED]	Storage Gateway does not have network connectivity.

Viewing your gateway system resource status

When your File Gateway starts, it checks its virtual CPU cores, root volume size, and RAM. It then determines whether the available system resources are sufficient for your gateway to function properly. You can view the results of the system resource check by using the gateway local console.

To view the status of a system resource check

1. Log in to the local console on your Amazon EC2 File Gateway. For instructions, see Logging in to your Amazon EC2 gateway local console.
2. From the **AWS Appliance Activation - Configuration** main menu, enter the corresponding numeral to select **View System Resource Check**.

   The gateway local console displays **[OK]**, **[WARNING]**, or **[FAIL]** to indicate the status of the resource as follows:

Message	Description
[OK]	The resource has passed the system resource check.
[WARNING]	The resource does not meet the recommend ed requirements, but your gateway can continue to function. The gateway local console displays a message that describes the results of the resource check.
[FAIL]	The resource does not meet the minimum requirements. Your gateway might not function properly. The gateway local console displays a message that describes the results of the resource check.

   The local console also displays the number of errors and warnings next to the resource check menu option.

Running Storage Gateway commands on the local console for an Amazon EC2 gateway

The AWS Storage Gateway console helps provide a secure environment for configuring and diagnosing issues with your gateway. Using the console commands, you can perform maintenance tasks such as saving routing tables or connecting to Support.

To run a configuration or diagnostic command

1. Log in to your gateway's local console. For instructions, see Logging in to your Amazon EC2 gateway local console.
2. From the **AWS Appliance Activation - Configuration** main menu, enter the corresponding numeral to select **Gateway Console**.
3. From the gateway console command prompt, enter **h**.

   The console displays the **AVAILABLE COMMANDS** menu, which lists the available commands:

Command	Function
dig	Collect output from dig for DNS troublesh ooting.
exit	Return to Configuration menu.
h	Display available command list.
ifconfig	View or configure network interfaces. **Note** We recommend configuring network or IP settings using the Storage Gateway console or the dedicated local console menu option. For instructions, see Configuring your gateway network settings.
ip	Show / manipulate routing, devices, and tunnels. **Note** We recommend configuring network or IP settings using the Storage Gateway console or the dedicated

Command	Function
	local console menu option. For instructions, see Configuring your gateway network settings.
iptables	Administration tool for IPv4 packet filtering and NAT.
ncport	Test connectivity to a specific TCP port on a network.
nping	Collect output from nping for network troubleshooting.
open-support-channel	Connect to AWS Support.
save-iptables	Persist IP tables.
save-routing-table	Save newly added routing table entry.
tcptraceroute	Collect traceroute output on TCP traffic to a destination.

4. From the gateway console command prompt, enter the corresponding command for the function you want to use, and follow the instructions.

To learn about a command, enter **man** + *command name* at the command prompt.

Configuring your Amazon EC2 gateway network settings

You can view and configure the network settings for your Amazon EC2 File Gateway by using the gateway local console.

To configure your network settings

1. Log in to the local console on your Amazon EC2 File Gateway. For instructions, see Logging in to your Amazon EC2 gateway local console.
2. From the **AWS Appliance Activation - Configuration** main menu, enter the corresponding numeral to select **Network Configuration**.

3. From the **AWS Appliance Activation - Network Configuration** menu, enter the corresponding numeral for the task that you want to perform:

   - **Edit DNS Configuration** - The gateway local console displays the available adapters for the primary and secondary DNS servers. The console then prompts you to provide the new IP address.
   - **View DNS Configuration** - The gateway local console displays the available adapters for the primary and secondary DNS servers.
   - **Configure Hostname** - The gateway local console prompts you to choose whether the gateway will use a static hostname that you specify, or if it will aquire a hostname automatically through DCHP or rDNS.

     > **Note**
     >
     > If you choose to configure a static hostname for your gateway, you must create an A record in your DNS system that points the IP address of the gateway to its static hostname.

   - **View Hostname Configuration** - The gateway local console displays hostname, aquisition mode, domain, and Active Directory realm for your Amazon EC2 File Gateway.

Shutting down your gateway VM

You might need to shutdown or reboot your VM for maintenance, such as when applying a patch to your hypervisor. You shut down on-premises gateway VMs using your hypervisor interface, and Amazon EC2 instances using the Amazon EC2 console.

> **Important**
>
> If you stop and start an Amazon EC2 gateway that uses ephemeral storage, the gateway will be permanently offline. This happens because the physical storage disk is replaced. There is no work-around for this issue. The only resolution is to delete the gateway and activate a new one on a new EC2 instance.

Replacing your existing FSx File Gateway with a new instance

You can replace an existing FSx File Gateway with a new instance as your data and performance needs grow, or if you receive an AWS notification to migrate your gateway. You might need to do this if you want to move your gateway to a better host platform or newer Amazon EC2 instances, or to refresh the underlying server hardware.

> **Note**
>
> Migration can only be performed between gateways of the same type. For example, you cannot migrate settings or data from an FSx File Gateway to an S3 File Gateway.

To replace your FSx File Gateway gateway with a new instance with an empty cache disk and a new Gateway ID:

1. Stop any applications that are writing to the existing FSx File Gateway. Verify that the `CachePercentDirty` metric on the **Monitoring** tab is `0` before you set up file system associations on the new gateway.
2. Use the AWS Command Line Interface (AWS CLI) to gather and save the configuration information about your existing FSx File Gateway and associated file systems by doing the following:

   a. Save the gateway configuration information for the FSx File Gateway.

   ```
   aws storagegateway describe-gateway-information --gateway-arn
    "arn:aws:storagegateway:us-east-2:123456789012:gateway/sgw-12A3456B"
   ```

   This command outputs a JSON block that contains metadata about the gateway, such as its name, network interfaces, configured time zone, and its state (whether the gateway is running).

   b. Save the Server Message Block (SMB) settings of the FSx File Gateway.

   ```
   aws storagegateway describe-smb-settings --gateway-arn
    "arn:aws:storagegateway:us-east-2:123456789012:gateway/sgw-12A3456B"
   ```

   This command outputs a JSON block that contains the domain name of the Microsoft Active Directory that the gateway is joined to.

c. Save file share information for each file system associated with the FSx File Gateway:

Use the following command for each associated file system.

```
aws storagegateway describe-file-system-associations --file-system-
association-arn-list "arn:aws:storagegateway:us-east-2:123456789012:fs-
association/fsa-987A654B"
```

This command outputs a JSON block that contains metadata about the file system, such as its location ARN, audit log destination, cache refresh attributes, configured IP addresses, and tags.

3. Create a new FSx File Gateway with the same settings and configuration as the old gateway. If necessary, refer to the information you saved in Step 2.
4. Create new file system associations for the new gateway with the same settings and configuration as the file systems that were configured on the old gateway. If necessary, refer to the information you saved in Step 2.
5. Confirm that your new gateway is working correctly, then remap/cut-over your clients from the old file systems to the new file systems in the manner that best suits your environment.
6. Confirm that your new gateway is working correctly, then delete the old gateway from the Storage Gateway console.

> **Important**
>
> Before you delete an FSx File Gateway, make sure that there are no applications currently writing to that gateway's cache. If you delete a gateway while it is in use, data loss can occur.

> **Warning**
>
> When a gateway is deleted, there is no way to recover it.

7. Delete the old gateway VM or Amazon EC2 instance.

Deleting your gateway and removing associated resources

If you don't plan to continue using your gateway, consider deleting the gateway and its associated resources. Removing resources avoids incurring charges for resources you don't plan to continue using and helps reduce your monthly bill.

When you delete a gateway, it no longer appears on the AWS Storage Gateway Management Console and its file system connections are closed. The procedure for deleting a gateway is the same for all gateway types; however, depending on the type of gateway you want to delete and the host it is deployed on, you follow specific instructions to remove associated resources.

You can delete a gateway using the Storage Gateway console or programmatically. You can find information following about how to delete a gateway using the Storage Gateway console. If you want to programmatically delete your gateway, see *AWS Storage Gateway API Reference*.

Deleting Your Gateway by Using the Storage Gateway Console

The procedure for deleting a gateway is the same for all gateway types. However, depending on the type of gateway you want to delete and the host the gateway is deployed on, you might have to perform additional tasks to remove resources associated with the gateway. Removing these resources helps you avoid paying for resources you don't plan to use.

> **Note**
>
> For gateways deployed on an Amazon EC2 instance, the instance continues to exist until you delete it.
> For gateways deployed on a virtual machine (VM), after you delete your gateway the gateway VM still exists in your virtualization environment. To remove the VM, use the VMware vSphere client, Microsoft Hyper-V Manager, or Linux Kernel-based Virtual Machine (KVM) client to connect to the host and remove the VM. Note that you can't reuse the deleted gateway's VM to activate a new gateway.

To delete a gateway

1. Open the Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. Choose **Gateways**, then select one or more gateways to delete.

3. For **Actions**, choose **Delete gateway**. The confirmation dialog box appears.

   > **Warning**
   >
   > Before you do this step, make sure that there are no applications currently writing to the gateway's volumes. If you delete the gateway while it is in use, data loss can occur. When a gateway is deleted, there is no way to get it back.

4. Verify that you want to delete the specified gateways, then type the word *delete* in the confirmation box, and choose **Delete**.
5. (Optional) If you want to provide feedback about your deleted gateway, complete the feedback dialog box, then choose **Submit**. Otherwise, choose **Skip**.

> **Important**
>
> You no longer pay software charges after you delete a gateway, but resources such as Amazon S3 bucket and Amazon EC2 instances persist. You can remove the gateway Amazon EC2 instance after the file gateway is removed.

Performance and optimization

This section describes guidance and best practices for optimizing File Gateway performance.

Topics

Basic performance guidance for FSx File Gateway

In this section, you can find guidance for provisioning hardware for your FSx File Gateway VM. The instance configurations that are listed in the table are examples, and are provided for reference.

For best performance, the cache disk size must be tuned to the size of the active working set. Using multiple local disks for the cache increases write performance by parallelizing access to data and leads to higher IOPS.

> **Note**
>
> We don't recommend using ephemeral storage. For information about using ephemeral storage, see Using ephemeral storage with EC2 gateways.
> The suggested size limit for individual directories in the file systems that you connect to File Gateway is 10,000 files per directory. You can use File Gateway with directories that have more than 10,000 files, but performance might be impacted.

In the following tables, *cache hit* read operations are reads from the file data that is served from cache. *Cache miss* read operations are reads from the file data that is served from Amazon FSx for Windows File Server.

The following table shows an example FSx File Gateway configuration.

FSx File Gateway performance on Windows clients

Example Configuration	Protocol	Write throughput (file sizes 1 GB)	Cache hit read throughput	Cache miss read throughput
Root disk: 80 GB, io1 SSD, 4,000 IOPS Cache disks: 2 x 2 TiB NVME Minimum network performance: 10 Gbps CPU: 32 vCPU | RAM: 244 GB	SMBv3 - 1 thread	162 MiB/sec (1.4 Gbps)	403 MiB/sec (3.4 Gbps)	288 MiB/sec (2.4 Gbps)
	SMBv3 - 8 threads	511 MiB/sec (4.3 Gbps)	571 MiB/sec (4.8 Gbps)	567 MiB/sec (4.8 Gbps)

Note

Your performance might vary based on your host platform configuration and network bandwidth. Write throughput performance decreases with file size, with the highest achievable throughput for small files (less than 32MiB) being 16 files per second.

Optimizing gateway performance

You can find information following about how to optimize the performance of your gateway. The guidance is based on adding resources to your gateway and adding resources to your application server.

Add Resources to Your Gateway

You can optimize gateway performance by adding resources to your gateway in one or more of the following ways.

Use higher-performance disks

To optimize gateway performance, you can add high-performance disks such as solid-state drives (SSDs) and a NVMe controller. You can also attach virtual disks to your VM directly from a storage area network (SAN) instead of the Microsoft Hyper-V NTFS. Improved disk performance generally results in better throughput and more input/output operations per second (IOPS). For information about adding disks, see Configuring additional cache storage.

To measure throughput, use the `ReadBytes` and `WriteBytes` metrics with the `Samples` Amazon CloudWatch statistic. For example, the `Samples` statistic of the `ReadBytes` metric over a sample period of 5 minutes divided by 300 seconds gives you the IOPS. As a general rule, when you review these metrics for a gateway, look for low throughput and low IOPS trends to indicate disk-related bottlenecks.

> **Note**
>
> CloudWatch metrics are not available for all gateways. For information about gateway metrics, see Monitoring your FSx File Gateway.

Add CPU resources to your gateway host

The minimum requirement for a gateway host server is four virtual processors. To optimize gateway performance, confirm that the four virtual processors that are assigned to the gateway VM are backed by four cores. In addition, confirm that you are not oversubscribing the CPUs of the host server.

When you add additional CPUs to your gateway host server, you increase the processing capability of the gateway. Doing this allows your gateway to deal with, in parallel, both storing data from your application to your local storage and uploading this data to FSx for Windows File Server. Additional CPUs also help ensure that your gateway gets enough CPU resources when the host is shared with other VMs. Providing enough CPU resources has the general effect of improving throughput.

Storage Gateway supports using 24 CPUs in your gateway host server. You can use 24 CPUs to significantly improve the performance of your gateway. We recommend the following gateway configuration for your gateway host server:

- 24 CPUs.
- 16 GiB of reserved RAM for File Gateways
  - 16 GiB of reserved RAM for gateways with cache size up to 16 TiB
  - 32 GiB of reserved RAM for gateways with cache size 16 TiB to 32 TiB
  - 48 GiB of reserved RAM for gateways with cache size 32 TiB to 64 TiB
- Disk 1 attached to paravirtual controller 1, to be used as the gateway cache as follows:
  - SSD using an NVMe controller.
- Network adapter 1 configured on VM network 1:
  - Use VM network 1 and add VMXnet3 (10 Gbps) to be used for ingestion.
- Network adapter 2 configured on VM network 2:
  - Use VM network 2 and add a VMXnet3 (10 Gbps) to be used to connect to AWS.

Back gateway virtual disks with separate physical disks

When you provision gateway disks, we strongly recommend that you don't provision local disks for local storage that use the same underlying physical storage disk. For example, for VMware ESXi, the underlying physical storage resources are represented as a data store. When you deploy the gateway VM, you choose a data store on which to store the VM files. When you provision a virtual disk (for example, as an upload buffer), you can store the virtual disk in the same data store as the VM or a different data store.

If you have more than one data store, then we strongly recommend that you choose one data store for each type of local storage you are creating. A data store that is backed by only one underlying physical disk can lead to poor performance. An example is when you use such a disk to back both the cache storage and upload buffer in a gateway setup. Similarly, a data store that is backed by a less high-performing RAID configuration such as RAID 1 can lead to poor performance.

Add Resources to Your Application Environment

Increase the bandwidth between your application server and your gateway

To optimize gateway performance, ensure that the network bandwidth between your application and the gateway can sustain your application needs. You can use the `ReadBytes` and `WriteBytes` metrics of the gateway to measure the total data throughput.

For your application, compare the measured throughput with the desired throughput. If the measured throughput is less than the desired throughput, then increasing the bandwidth between your application and gateway can improve performance if the network is the bottleneck. Similarly, you can increase the bandwidth between your VM and your local disks, if they're not direct-attached.

Add CPU resources to your application environment

If your application can use additional CPU resources, then adding more CPUs can help your application to scale its I/O load.

Some file operations on the FSx File Gateway, such as top-level folder renames or permission changes, can result in multiple file operations that lead to a high I/O load on your FSx for Windows File Server file system. If your file system doesn't have enough performance resources for your workload, the file system might delete shadow copies because it prioritizes availability for ongoing I/O over historical shadow copy retention.

In the Amazon FSx console, check the **Monitoring and performance** page to see if your file system is under-provisioned. If it is, you can switch to SSD storage, increase throughput capacity, or increase SSD IOPS to handle your workload.

Maximizing S3 File Gateway throughput

The following sections describe best practices for maximizing throughput between your NFS and SMB clients, S3 File Gateway, and Amazon S3. The guidance provided in each section contributes incrementally to improving overall throughput. While none of these recommendations are required, and they are not interdependent, they have been selected and ordered in a logical way that Support uses to test and tune S3 File Gateway implementations. As you implement and test these suggestions, keep in mind that each S3 File Gateway deployment is unique, so your results may vary.

S3 File Gateway provides a file interface to store and retrieve Amazon S3 objects using industry-standard NFS or SMB file protocols, with a native 1:1 mapping between file and object. You deploy S3 File Gateway as a virtual machine either on-premises in your VMware, Microsoft Hyper-V, or Linux KVM environment, or in the AWS cloud as an Amazon EC2 instance. S3 File Gateway is not designed to act as a full enterprise NAS replacement. S3 File Gateway emulates a file system, but it is not a file system. Using Amazon S3 as durable backend storage creates additional overhead on each I/O operation, so evaluating S3 File Gateway performance against an existing NAS or file server is not an equivalent comparison.

Deploy your gateway in the same location as your clients

We recommend deploying your S3 File Gateway virtual appliance in a physical location with as little network latency as possible between it and your NFS or SMB clients. When choosing a location for your gateway, consider the following:

- Lower network latency to the gateway can help improve performance of NFS or SMB clients.
- S3 File Gateway is designed to tolerate higher network latency between the gateway and Amazon S3 than between the gateway and the clients.
- For S3 File Gateway instances deployed in Amazon EC2, we recommend keeping the gateway and NFS or SMB clients in the same placement group. For more information, see Placement groups for your Amazon EC2 instances in the Amazon Elastic Compute Cloud User Guide.

Reduce bottlenecks caused by slow disks

We recommend monitoring the `IoWaitPercent` CloudWatch metric to identify performance bottlenecks that can result from slow storage disks on your S3 File Gateway. When attempting to optimize disk-related performance issues, consider the following:

- `IoWaitPercent` reports the percentage of time that the CPU is waiting for a response from the root or cache disks.
- When `IoWaitPercent` is greater than 5-10%, this usually indicates a gateway performance bottleneck caused by underperforming disks. This metric should be as close to 0% as possible - meaning that the gateway is never waiting on the disk - which helps to optimize CPU resources.
- You can check `IoWaitPercent` on the **Monitoring** tab of the Storage Gateway console, or configure recommended CloudWatch alarms to notify you automatically if the metric spikes above a specific threshold. For more information, see Creating recommended CloudWatch alarms for your gateway.

- We recommend using either NVMe or SSD for your gateway's root and cache disks to minimize `IoWaitPercent`.

Adjust virtual machine resource allocation for CPU, RAM, and cache disks

When attempting to optimize throughput for your S3 File Gateway, it is important to allocate sufficient resources to the gateway VM, including CPU, RAM, and cache disks. The minimum virtual resource requirements of 4 CPUs, 16GB RAM, and 150GB cache storage are typically only suitable for smaller workloads. When allocating virtual resources for larger workloads, we recommend the following:

- Increase the allocated number of CPUs to between 16 and 48, depending on the typical CPU usage generated by your S3 File Gateway. You can monitor CPU usage using the `UserCpuPercent` metric. For more information, see Understanding gateway metrics.
- Increase the allocated RAM to between 32 and 64 GB.

> **Note**
>
> S3 File Gateway cannot utilize more than 64 GB of RAM.

- Use NVMe or SSD for root disks and cache disk, and size your cache disks to align with the peak working data set that you plan to write to the gateway. For more information, see S3 File Gateway cache sizing best practices on the official Amazon Web Services YouTube channel.
- Add at least 4 virtual cache disks to the gateway, rather than using a single large disk. Multiple virtual disks can improve performance even if they share the same underlying physical disk, but improvements are typically greater when the virtual disks are located on different underlying physical disks.

  For example, if you want to deploy 12TB of cache, you could use one of the following configurations:

  - 4 x 3 TB cache disks
  - 8 x 1.5 TB cache disks
  - 12 x 1 TB cache disks

In addition to performance, this allows for more efficient management of the virtual machine over time. As your workload changes, you can incrementally increase the number of cache disks and your overall cache capacity, while maintaining the original size of each individual virtual disk to preserve gateway integrity.

For more information, see Deciding the amount of local disk storage.

When deploying S3 File Gateway as an Amazon EC2 instance, consider the following:

- The instance type you choose can significantly impact gateway performance. Amazon EC2 provides broad flexibility for adjusting the resource allocation for your S3 File Gateway instance.
- For recommended Amazon EC2 instance types for S3 File Gateway, see Requirements for Amazon EC2 instance types.
- You can change the Amazon EC2 instance type that hosts an active S3 File Gateway. This allows you to easily adjust the Amazon EC2 hardware generation and resource allocation to find an ideal price-to-performance ratio. To change the instance type, use the following procedure in the Amazon EC2 console:

  1. Stop the Amazon EC2 instance.
  2. Change the Amazon EC2 instance type.
  3. Power on the Amazon EC2 instance.

  > **Note**
  >
  > Stopping an instance that hosts an S3 File Gateway will temporarily disrupt file share access. Make sure to schedule a maintenance window if necessary.

- The price-to-performance ratio of an Amazon EC2 instance refers to how much computing power you get for the price you pay. Typically, newer generation Amazon EC2 instances offer the best price-to-performance ratio, with newer hardware and improved performance at a relatively lower cost compared to older generations. Factors such as instance type, region, and usage patterns impact this ratio, so it is important to select the right instance for your specific workload to optimize cost-effectiveness.

Adjust the SMB security level

The SMBv3 protocol allows for both SMB signing and SMB encryption, which have some trade-offs in performance and security. To optimize throughput, you can adjust your gateway's SMB security level to specify which of these security features are enforced for client connections. For more information, see Set a security level for your gateway.

When adjusting the SMB security level, consider the following:

- The default security level for S3 File Gateway is **Enforce encryption**. This setting enforces both encryption and signing for SMB client connections to gateway file shares, meaning that all traffic from the client to the gateway is encrypted. This setting does not affect traffic from the gateway to AWS, which is always encrypted.

  The gateway limits each encrypted client connection to a single vCPU. For example, if you have only 1 encrypted client, then that client will be limited to only 1 vCPU, even if 4 or more vCPUs are allocated to the gateway. Because of this, throughput for encrypted connections from a single client to S3 File Gateway is typically bottlenecked between 40-60 MB/s.
- If your security requirements allow for a more relaxed posture, you can change the security level to **Client negotiated**, which will disable SMB encryption and enforce SMB signing only. With this setting, client connections to the gateway can utilize multiple vCPUs, which typically results in increased throughput performance.

> **Note**
>
> After you change the SMB security level for your S3 File Gateway, you must wait for the file share status to change from **Updating** to **Available** in the Storage Gateway console, and then disconnect and reconnect your SMB clients for the new setting to take effect.

Use multiple threads and clients to parallelize write operations

It is difficult to achieve maximum throughput performance with an S3 File Gateway that uses only one NFS or SMB client to write one file at a time, because sequential writing from a single client is a single-threaded operation. Instead, we recommend using multiple threads from each NFS or SMB client to write multiple files in parallel, and using multiple NFS or SMB clients simultaneously to your S3 File Gateway to maximize the gateway throughput.

Using multiple threads can significantly improve performance. However, using more threads requires more system resources, which can negatively impact performance if the gateway is not sized to meet the increased load. In a typical deployment, you can expect to achieve better throughput performance as you add more threads and clients, until you reach the maximum hardware and bandwidth limitations for your gateway. We recommend experimenting with different thread counts to find the optimal balance between speed and system resource usage for your specific hardware and network configuration.

Consider the following information about common tools that can help you test your thread and client configuration:

- You can test multithreaded write performance by using tools such as robocopy to copy a set of files to a file share on your gateway. By default, robocopy uses 8 threads when copying files, but you can specify up to 128 threads.

  To use multiple threads with robocopy, add the `/MT:n` switch to your command, where `n` is the number of threads you want to use. For example:

  ```
  robocopy C:\source D:\destination /MT:64
  ```

  This command will use 64 threads for the copy operation.

  > **Note**
  >
  > We don't recommend using Windows Explorer to drag and drop files when testing for maximum throughput, as this method is limited to a single thread and copies the files sequentially.

  For more information, see robocopy on the Microsoft Learn website.
- You can also conduct tests using common storage benchmarking tools such as DISKSPD, or FIO. These tools have options to adjust the number of threads, I/O depth, and other parameters to match your specific workload requirements.

  DiskSpd allows you to control the number of threads using the `-t` parameter. For example:

  ```
  diskspd -c10G -d300 -r -w50 -t64 -o32 -b1M -h -L C:\testfile.dat
  ```

  This example command does the following:

- Creates a 10GB test file (`-c1G`)
- Runs for 300 seconds (`-d300`)
- Performs random I/O test with 50% reads 50% writes (`-r -w50`)
- Uses 64 threads (`-t64`)
- Sets queue depth to 32 per thread (`-o32`)
- Uses 1MB block size (`-b1M`)
- Disables hardware and software caching (`-h -L`)

For more information, see Use DISKSPD to test workload storage performance on the Microsoft Learn website.

- FIO uses the `numjobs` parameter to control the number of parallel threads. For example:

```
fio --name=mixed_test --rw=randrw --rwmixread=70 --bs=1M -- iodepth=64
--size=10G --runtime=300 --numjobs=64 --ioengine=libaio --direct=1 --
group_reporting
```

This example command does the following:

- Performs random I/O test (`--rw=randrw`)
- Performs 70% reads and 30% writes (`--rwmixread=70`)
- Uses 1MB block size (`--bs=1M`)
- Sets I/O depth to 64 (`--iodepth=64`)
- Tests on a 10 GB file (`--size=10G`)
- Runs for 5 minutes (`--runtime=300`)
- Creates 64 parallel jobs (threads) (`--numjobs=64`)
- Uses asynchronous I/O engine (`--ioengine=libaio`)
- Groups results for easier analysis (`--group_reporting`)

For more information, see the fio Linux man page.

-

Turn off automated cache refresh

The automated cache refresh feature allows your S3 File Gateway to refresh its metadata automatically, which can help capture any changes that users or applications make to your file

set by writing to the Amazon S3 bucket directly, rather than through the gateway. For more information, see Refreshing Amazon S3 bucket object cache.

To optimize gateway throughput, we recommend turning this feature off in deployments where all reads and writes to the Amazon S3 bucket will be performed through your S3 File Gateway.

When configuring automated cache refresh, consider the following:

- If you need to use automated cache refresh because users or applications in your deployment do occasionally write to Amazon S3 directly, then we recommend configuring the longest possible time interval between refreshes that is still practical for your business needs. A longer cache refresh interval helps reduce the number of metadata operations that the gateway needs to perform when browsing directories or modifying files.

  For example: set automated cache refresh to 24 hours, rather than 5 minutes, if that is tolerable for your workload.
- The minimum time interval is 5 minutes. The maximum interval is 30 days.
- If you choose to set a very short cache refresh interval, we recommend testing the directory browsing experience for your NFS and SMB clients. The time it takes to refresh the gateway cache can increase substantially depending on the number of files and subdirectories in your Amazon S3 bucket.

Increase the number of Amazon S3 uploader threads

By default, S3 File Gateway opens 8 threads for Amazon S3 data upload, which provides sufficient upload capacity for most typical deployments. However, it is possible for a gateway to receive data from NFS and SMB clients at a higher rate than it can upload to Amazon S3 with the standard 8 thread capacity, which can cause the local cache to reach its storage limit.

In specific circumstances, Support can increase the Amazon S3 upload thread pool count for your gateway from 8 to 40, which allows more data to be uploaded in parallel. Depending on bandwidth and other factors specific to your deployment, this can significantly increase upload performance and help reduce the amount of cache storage needed to support your workload.

We recommend using the `CachePercentDirty` CloudWatch metric to monitor the amount of data stored on the local gateway cache disks that has not yet been uploaded to Amazon S3, and contacting Support to help determine if increasing the upload thread pool count might improve throughput for your S3 File Gateway. For more information, see Understanding gateway metrics.

> **Note**
>
> This setting consumes additional gateway CPU resources. We recommend monitoring gateway CPU usage and increasing allocated CPU resources if necessary.

Increase SMB timeout settings

When S3 File Gateway copies large files to an SMB file share, the SMB client connection can timeout after an extended period of time.

We recommend extending the SMB session timeout setting for your SMB clients to 20 minutes or more, depending on the size of the files and the write speed of your gateway. The default is 300 seconds, or 5 minutes. For more information, see Your gateway backup job fails or there are errors when writing to your gateway.

Turn on opportunistic locking for compatible applications

Opportunistic locking, or "oplocks", is enabled by default for each new S3 File Gateway. When using oplocks with compatible applications, the client batches multiple smaller operations into larger ones, which is more efficient for the client, the gateway, and the network. We recommend keeping opportunistic locking turned on if you use applications that leverage client-side local caching, such as Microsoft Office, Adobe Suite, and many others, because it can significanty improve performance.

If you turn opportunistic locking off, applications that support oplocks will typically open large files (50 MB or larger) much more slowly. This delay occurs because the gateway sends data in 4 KB parts, which results in high I/O and low throughput.

Adjust gateway capacity according to the size of the working file set

The gateway capacity parameter specifies the maximum number of files for which your gateway will store metadata in its local cache. By default, gateway capacity is set to **Small**, which means the gateway stores metadata for up to 5 million files. The default setting works well for most workloads, even if there are hundreds of millions, or even billions of objects in Amazon S3, because only a small subset of files are actively accessed at a given time in a typical deployment. This group of files is referred to as the "working set".

If your workload regularly accesses a working set of files greater than 5 million, then your gateway will need to perform frequent cache evictions, which are small I/O operations that are stored in RAM and persisted on the root disk. This can negatively impact gateway performance as the gateway fetches fresh data from Amazon S3.

You can monitor the `IndexEvictions` metric to determine the number of files whose metadata was evicted from the cache to make room for new entries. For more information, see Understanding gateway metrics.

We recommend using the `UpdateGatewayInformation` API action to increase the gateway capacity to correspond with the number of files in your typical working set. For more information, see UpdateGatewayInformation.

Note

Increasing the gateway capacity requires additional RAM and root disk capacity.

- Small (5 million files) requires at least 16 GB of RAM and 80 GB root disk.
- Medium (10 million files) requires at least 32 GB of RAM and 160 GB root disk.
- Large (20 million files) requires 64 GB of RAM and 240 GB root disk.

Important

Gateway capacity cannot be decreased.

Deploy multiple gateways for larger workloads

We recommend splitting your workload across multiple gateways when possible, rather than consolidating many file shares on a single large gateway. For example, you could isolate one heavily-used file share on one gateway, while grouping the less frequently used file shares together on another gateway.

When planning a deployment with multiple gateways and file shares, consider the following:

- The maximum number of file shares on a single gateway is 50, but the number of file shares managed by a gateway can impact the gateway's performance. For more information, see Performance guidance for gateways with multiple file shares.

- Resources on each S3 File Gateway are shared across all file shares, without partitioning.
- A single file share with heavy usage can impact the performance of other file shares on the gateway.

> **Note**
>
> We do not recommended creating multiple file shares that are mapped to the same Amazon S3 location from multiple gateways, unless at least one of them is read-only. Simultaneous writes to the same file from multiple gateways is considered a multi-writer scenario, which can cause data integrity issues.

Optimizing S3 File Gateway for SQL Server database backups

Database backups are a common and recommended use case for S3 File Gateway, which provides cost-effective short and long term retention by storing database backups in Amazon S3, with the ability to lifecycle to lower cost storage tiers as needed. With this solution, you can reduce the need for enterprise backup applications using built-in tools such as SQL Server Management Studio and Oracle RMAN.

The following sections describe best practices to tune your S3 File Gateway deployment for optimized performance and cost-effective support for hundreds of terabytes of SQL database backups. The guidance provided in each section contributes incrementally to improving overall throughput. While none of these recommendations are required, and they are not interdependent, they have been selected and ordered in a logical way that Support uses to test and tune S3 File Gateway implementations. As you implement and test these suggestions, keep in mind that each S3 File Gateway deployment is unique, so your results may vary.

S3 File Gateway provides a file interface to store and retrieve Amazon S3 objects using industry-standard NFS or SMB file protocols, with a native 1:1 mapping between file and object. You deploy S3 File Gateway as a virtual machine either on-premises in your VMware, Microsoft Hyper-V, or Linux KVM environment, or in the AWS cloud as an Amazon EC2 instance. S3 File Gateway is not designed to act as a full enterprise NAS replacement. S3 File Gateway emulates a file system, but it is not a file system. Using Amazon S3 as durable backend storage creates additional overhead on each I/O operation, so evaluating S3 File Gateway performance against an existing NAS or file server is not an equivalent comparison.

Deploy your gateway in the same location as your SQL Servers

We recommend deploying your S3 File Gateway virtual appliance in a physical location with as little network latency as possible between it and your SQL servers. When choosing a location for your gateway, consider the following:

- Lower network latency to the gateway can help improve performance of SMB clients, such as SQL servers.
- S3 File Gateway is designed to tolerate higher network latency between the gateway and Amazon S3 than between the gateway and the clients.
- For S3 File Gateway instances deployed in Amazon EC2, we recommend keeping the gateway and SQL servers in the same placement group. For more information, see Placement groups for your Amazon EC2 instances in the Amazon Elastic Compute Cloud User Guide.

Reduce bottlenecks caused by slow disks

We recommend monitoring the `IoWaitPercent` CloudWatch metric to identify performance bottlenecks that can result from slow storage disks on your S3 File Gateway. When attempting to optimize disk-related performance issues, consider the following:

- `IoWaitPercent` reports the percentage of time that the CPU is waiting for a response from the root or cache disks.
- When `IoWaitPercent` is greater than 5-10%, this usually indicates a gateway performance bottleneck caused by underperforming disks. This metric should be as close to 0% as possible - meaning that the gateway is never waiting on the disk - which helps to optimize CPU resources.
- You can check `IoWaitPercent` on the **Monitoring** tab of the Storage Gateway console, or configure recommended CloudWatch alarms to notify you automatically if the metric spikes above a specific threshold. For more information, see Creating recommended CloudWatch alarms for your gateway.
- We recommend using either NVMe or SSD for your gateway's root and cache disks to minimize `IoWaitPercent`.

Adjust S3 File Gateway virtual machine resource allocation for CPU, RAM, and cache disks

When attempting to optimize throughput for your S3 File Gateway, it is important to allocate sufficient resources to the gateway VM, including CPU, RAM, and cache disks. The minimum virtual resource requirements of 4 CPUs, 16GB RAM, and 150GB cache storage are typically only suitable for smaller workloads. When allocating virtual resources for larger workloads, we recommend the following:

- Increase the allocated number of CPUs to between 16 and 48, depending on the typical CPU usage generated by your S3 File Gateway. You can monitor CPU usage using the `UserCpuPercent` metric. For more information, see Understanding gateway metrics.
- Increase the allocated RAM to between 32 and 64 GB.

  > **Note**
  >
  > S3 File Gateway cannot utilize more than 64 GB of RAM.

- Use NVMe or SSD for root disks and cache disk, and size your cache disks to align with the peak working data set that you plan to write to the gateway. For more information, see S3 File Gateway cache sizing best practices on the official Amazon Web Services YouTube channel.
- Add at least 4 virtual cache disks to the gateway, rather than using a single large disk. Multiple virtual disks can improve performance even if they share the same underlying physical disk, but improvements are typically greater when the virtual disks are located on different underlying physical disks.

  For example, if you want to deploy 12TB of cache, you could use one of the following configurations:

  - 4 x 3 TB cache disks
  - 8 x 1.5 TB cache disks
  - 12 x 1 TB cache disks

  In addition to performance, this allows for more efficient management of the virtual machine over time. As your workload changes, you can incrementally increase the number of cache disks and your overall cache capacity, while maintaining the original size of each individual virtual disk to preserve gateway integrity.

For more information, see Deciding the amount of local disk storage.

When deploying S3 File Gateway as an Amazon EC2 instance, consider the following:

- The instance type you choose can significantly impact gateway performance. Amazon EC2 provides broad flexibility for adjusting the resource allocation for your S3 File Gateway instance.
- For recommended Amazon EC2 instance types for S3 File Gateway, see Requirements for Amazon EC2 instance types.
- You can change the Amazon EC2 instance type that hosts an active S3 File Gateway. This allows you to easily adjust the Amazon EC2 hardware generation and resource allocation to find an ideal price-to-performance ratio. To change the instance type, use the following procedure in the Amazon EC2 console:

  1. Stop the Amazon EC2 instance.
  2. Change the Amazon EC2 instance type.
  3. Power on the Amazon EC2 instance.

  > **Note**
  >
  > Stopping an instance that hosts an S3 File Gateway will temporarily disrupt file share access. Make sure to schedule a maintenance window if necessary.

- The price-to-performance ratio of an Amazon EC2 instance refers to how much computing power you get for the price you pay. Typically, newer generation Amazon EC2 instances offer the best price-to-performance ratio, with newer hardware and improved performance at a relatively lower cost compared to older generations. Factors such as instance type, region, and usage patterns impact this ratio, so it is important to select the right instance for your specific workload to optimize cost-effectiveness.

Improve SMB client throughput by adjusting the security level of your S3 File Gateway

The SMBv3 protocol allows for both SMB signing and SMB encryption, which have some trade-offs in performance and security. To optimize throughput, you can adjust your gateway's SMB security level to specify which of these security features are enforced for client connections. For more information, see Set a security level for your gateway.

When adjusting the SMB security level, consider the following:

- The default security level for S3 File Gateway is **Enforce encryption**. This setting enforces both encryption and signing for SMB client connections to gateway file shares, meaning that all traffic from the client to the gateway is encrypted. This setting does not affect traffic from the gateway to AWS, which is always encrypted.

  The gateway limits each encrypted client connection to a single vCPU. For example, if you have only 1 encrypted client, then that client will be limited to only 1 vCPU, even if 4 or more vCPUs are allocated to the gateway. Because of this, throughput for encrypted connections from a single client to S3 File Gateway is typically bottlenecked between 40-60 MB/s.
- If your security requirements allow for a more relaxed posture, you can change the security level to **Client negotiated**, which will disable SMB encryption and enforce SMB signing only. With this setting, client connections to the gateway can utilize multiple vCPUs, which typically results in increased throughput performance.

> **Note**
>
> After you change the SMB security level for your S3 File Gateway, you must wait for the file share status to change from **Updating** to **Available** in the Storage Gateway console, and then disconnect and reconnect your SMB clients for the new setting to take effect.

Improve SMB client throughput by splitting SQL backups into multiple files

- It is difficult to achieve the maximum throughput performance with an S3 File Gateway that only one SQL server writing one file at a time, because sequential writing from a single SQL server is a single-threaded operation. Instead, we recommend using multiple threads from each SQL server to write multiple files in parallel, and using multiple SQL servers simultaneously to your S3 File Gateway to maximize the gateway throughput. With SQL backups, splitting backups into multiple files allows each file to utilize a separate thread, which will write multiple files simultaneously to the S3 File Gateway file share. The more threads you have, the more throughput you can achieve, up to the limits of the gateway.
- SQL Server supports writing to multiple files at the same time during a single backup operation. For instance, you can specify multiple file destinations using T-SQL commands or SQL Server Management Studio (SSMS). Each file uses a separate thread to send data from the SQL server

to the gateway file share. This approach allows for better I/O throughput, which can significantly improve backup speed and efficiency.

When configuring your SQL server backups, consider the following:

- By splitting backups into multiple files, SQL Server admins can optimize backup times and manage large database backups more effectively.
- The number of files used depends on the server's storage configuration and performance requirements. For large databases, we recommend breaking backups into several smaller files between 10 GB and 20 GB each.
- There is no strict limit on how many files SQL Server can write to during a backup, but practical considerations like storage architecture and network bandwidth should guide this choice.

For more information, see:

- Back up SQL Server 43-67% faster by writing to multiple files
- Easily store your SQL Server backups in Amazon S3 using File Gateway

Prevent large file copy failures by increasing SMB timeout settings

When S3 File Gateway copies large SQL backup files to an SMB file share, the SMB client connection can timeout after an extended period of time. We recommend extending the SMB session timeout setting for your SQL server SMB clients to 20 minutes or more, depending on the size of the files and the write speed of your gateway. The default is 300 seconds, or 5 minutes. For more information, see Your gateway backup job fails or there are errors when writing to your gateway.

Increase the number of Amazon S3 uploader threads

By default, S3 File Gateway opens 8 threads for Amazon S3 data upload, which provides sufficient upload capacity for most typical deployments. However, it is possible for a gateway to receive data from SQL servers at a higher rate than it can upload to Amazon S3 with the standard 8 thread capacity, which can cause the local cache to reach its storage limit.

In specific circumstances, Support can increase the Amazon S3 upload thread pool count for your gateway from 8 to 40, which allows more data to be uploaded in parallel. Depending on bandwidth

and other factors specific to your deployment, this can significantly increase upload performance and help reduce the amount of cache storage needed to support your workload.

We recommend using the `CachePercentDirty` CloudWatch metric to monitor the amount of data stored on the local gateway cache disks that has not yet been uploaded to Amazon S3, and contacting Support to help determine if increasing the upload thread pool count might improve throughput for your S3 File Gateway. For more information, see Understanding gateway metrics.

> **Note**
>
> This setting consumes additional gateway CPU resources. We recommend monitoring gateway CPU usage and increasing allocated CPU resources if necessary.

Turn off automated cache refresh

The automated cache refresh feature allows your S3 File Gateway to refresh its metadata automatically, which can help capture any changes that users or applications make to your file set by writing to the Amazon S3 bucket directly, rather than through the gateway. For more information, see Refreshing Amazon S3 bucket object cache.

To optimize gateway throughput, we recommend turning this feature off in deployments where all reads and writes to the Amazon S3 bucket will be performed through your S3 File Gateway.

When configuring automated cache refresh, consider the following:

- If you need to use automated cache refresh because users or applications in your deployment do occasionally write to Amazon S3 directly, then we recommend configuring the longest possible time interval between refreshes that is still practical for your business needs. A longer cache refresh interval helps reduce the number of metadata operations that the gateway needs to perform when browsing directories or modifying files.

  For example: set automated cache refresh to 24 hours, rather than 5 minutes, if that is tolerable for your workload.
- The minimum time interval is 5 minutes. The maximum interval is 30 days.
- If you choose to set a very short cache refresh interval, we recommend testing the directory browsing experience for your SQL servers. The time it takes to refresh the gateway cache can increase substantially depending on the number of files and subdirectories in your Amazon S3 bucket.

Deploy multiple gateways to support the workload

It is possible for Storage Gateway to support SQL backups for large environments with hundreds of SQL databases, multiple SQL Servers, and hundreds of terabytes of backup data by splitting the workload across multiple gateways.

When planning a deployment with multiple gateways and SQL servers, consider the following:

- A single gateway can typically upload up to 20 TB per day, with sufficient hardware resources and bandwidth. You can increase this limit up to 40 TB per day by increasing the number of Amazon S3 uploader threads.
- We recommend conducting a proof-of-concept test to measure performance and account for all of the variables in your deployment. After you determine the peak throughput of your SQL backup workload, you can scale the number of gateways to meet your requirements.
- We recommend designing your solution with growth in mind, because the number of databases and size of databases can increase over time. To continue to scale and support an increasing workload, you can deploy additional gateways as needed.

Additional resources for database backup workloads

- Store SQL Server backups in Amazon S3 using AWS Storage Gateway
- Easily store your SQL Server backups in Amazon S3 using File Gateway
- Using AWS Storage Gateway to store Oracle database backups in Amazon S3
- Backing up Oracle databases to Amazon S3 at scale
- Integrate an SAP ASE database to Amazon S3 using AWS Storage Gateway
- How one AWS Hero uses AWS Storage Gateway for in-cloud backup
- S3 File Gateway cache sizing best practices

Security in AWS Storage Gateway

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from a data center and network architecture that is built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. The shared responsibility model describes this as security *of* the cloud and security *in* the cloud:

- **Security of the cloud** – AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS 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 AWS Compliance Programs. To learn about the compliance programs that apply to AWS Storage Gateway, see AWS Services in Scope by Compliance Program.
- **Security in the cloud** – Your responsibility is determined by the AWS 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 Storage Gateway. The following topics show you how to configure Storage Gateway to meet your security and compliance objectives. You also learn how to use other AWS services that help you to monitor and secure your Storage Gateway resources.

Data protection in AWS Storage Gateway

The AWS shared responsibility model applies to data protection in AWS Storage Gateway. As described in this model, AWS is responsible for protecting the global infrastructure that runs all of the AWS 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 AWS services that you use. For more information about data privacy, see the Data Privacy FAQ. For information about data protection in Europe, see the AWS Shared Responsibility Model and GDPR blog post on the *AWS Security Blog*.

For data protection purposes, we recommend that you protect AWS account credentials and set up individual users with AWS IAM Identity Center or AWS 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 AWS resources. We require TLS 1.2 and recommend TLS 1.3.
- Set up API and user activity logging with AWS CloudTrail. For information about using CloudTrail trails to capture AWS activities, see Working with CloudTrail trails in the *AWS CloudTrail User Guide*.
- Use AWS encryption solutions, along with all default security controls within AWS 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 140-3 validated cryptographic modules when accessing AWS 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) 140-3.

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 Storage Gateway or other AWS services using the console, API, AWS CLI, or AWS 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 using AWS KMS

Amazon FSx File Gateway supports SMB encryption up to the latest SMB v3.1.1 specification, including AES 128 CCM and AES 128 GCM. Compatible clients will connect using encryption automatically. Additionally, FSx File Gateway uses SMB encryption when it communicates with FSx for Windows File Server in AWS. You must configure an AWS Direct Connect link to AWS, and set appropriate policies to allow SMB traffic and management traffic to pass through to AWS.

Encrypting a file system

For information see, Data Encryption in Amazon FSx in the *Amazon FSx for Windows File Server User Guide*.

When using AWS KMS to encrypt your data, keep the following in mind:

- Your data is encrypted at rest in the cloud. That is, the data is encrypted in Amazon FSx.

- IAM users must have the required permissions to call the AWS KMS API operations. For more information, see Using IAM policies with AWS KMS in the *AWS Key Management Service Developer Guide*.

> **Important**
>
> When you use an AWS KMS key for server-side encryption, you must choose a symmetric key. Storage Gateway does not support asymmetric keys. For more information, see Using symmetric and asymmetric keys in the *AWS Key Management Service Developer Guide*.

For more information about AWS KMS, see What is AWS Key Management Service?

Identity and access management for AWS Storage Gateway

AWS Identity and Access Management (IAM) is an AWS service that helps an administrator securely control access to AWS resources. IAM administrators control who can be *authenticated* (signed in) and *authorized* (have permissions) to use AWS SGW resources. IAM is an AWS service that you can use with no additional charge.

Topics

Audience

How you use AWS Identity and Access Management (IAM) differs, depending on the work that you do in AWS SGW.

Service user – If you use the AWS SGW service to do your job, then your administrator provides you with the credentials and permissions that you need. As you use more AWS SGW 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 AWS SGW, see Troubleshooting AWS Storage Gateway identity and access.

Service administrator – If you're in charge of AWS SGW resources at your company, you probably have full access to AWS SGW. It's your job to determine which AWS SGW 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 AWS SGW, see How AWS Storage Gateway 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 AWS SGW. To view example AWS SGW identity-based policies that you can use in IAM, see Identity-based policy examples for AWS Storage Gateway.

Authenticating with identities

Authentication is how you sign in to AWS using your identity credentials. You must be *authenticated* (signed in to AWS) as the AWS account root user, as an IAM user, or by assuming an IAM role.

You can sign in to AWS as a federated identity by using credentials provided through an identity source. AWS IAM Identity Center (IAM Identity Center) users, your company's single sign-on authentication, and your Google or Facebook credentials are examples of federated identities. When you sign in as a federated identity, your administrator previously set up identity federation using IAM roles. When you access AWS by using federation, you are indirectly assuming a role.

Depending on the type of user you are, you can sign in to the AWS Management Console or the AWS access portal. For more information about signing in to AWS, see How to sign in to your AWS account in the *AWS Sign-In User Guide*.

If you access AWS programmatically, AWS 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 AWS tools, you must sign requests yourself. For more information about using the recommended method to sign requests yourself, see AWS Signature Version 4 for 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, AWS recommends that you use multi-factor authentication (MFA) to increase the security of your account. To learn more, see Multi-factor authentication in

the *AWS IAM Identity Center User Guide* and AWS Multi-factor authentication in IAM in the *IAM User Guide*.

AWS account root user

When you create an AWS account, you begin with one sign-in identity that has complete access to all AWS services and resources in the account. This identity is called the AWS 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 AWS services by using temporary credentials.

A *federated identity* is a user from your enterprise user directory, a web identity provider, the AWS Directory Service, the Identity Center directory, or any user that accesses AWS services by using credentials provided through an identity source. When federated identities access AWS accounts, they assume roles, and the roles provide temporary credentials.

For centralized access management, we recommend that you use AWS IAM Identity Center. You can create users and groups in IAM Identity Center, or you can connect and synchronize to a set of users and groups in your own identity source for use across all your AWS accounts and applications. For information about IAM Identity Center, see What is IAM Identity Center? in the *AWS IAM Identity Center User Guide*.

IAM users and groups

An *IAM user* is an identity within your AWS 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 Use cases for IAM users in the *IAM User Guide*.

IAM roles

An *IAM role* is an identity within your AWS account that has specific permissions. It is similar to an IAM user, but is not associated with a specific person. To temporarily assume an IAM role in the AWS Management Console, you can switch from a user to an IAM role (console). You can assume a role by calling an AWS CLI or AWS API operation or by using a custom URL. For more information about methods for using roles, see Methods to assume a role 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 Create a role for a third-party identity provider (federation) in the *IAM User Guide*. If you use IAM Identity Center, you configure a permission set. To control what your identities can access after they authenticate, IAM Identity Center correlates the permission set to a role in IAM. For information about permissions sets, see Permission sets in the *AWS IAM Identity Center 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 AWS 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 Cross account resource access in IAM in the *IAM User Guide*.
- **Cross-service access** – Some AWS services use features in other AWS 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 AWS, 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 AWS service, combined with the requesting AWS service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other AWS 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 Create a role to delegate permissions to an AWS service in the *IAM User Guide*.
- **Service-linked role** – A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your AWS 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 AWS CLI or AWS API requests. This is preferable to storing access keys within the EC2 instance. To assign an AWS 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 Use an IAM role to grant permissions to applications running on Amazon EC2 instances in the *IAM User Guide*.

Managing access using policies

You control access in AWS by creating policies and attaching them to AWS identities or resources. A policy is an object in AWS that, when associated with an identity or resource, defines their permissions. AWS 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 AWS 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 AWS 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 AWS Management Console, the AWS CLI, or the AWS 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 Define custom IAM permissions with customer managed 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 AWS account. Managed policies include AWS managed policies and customer managed policies. To learn how to choose between a managed policy or an inline policy, see Choose 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 AWS services.

Resource-based policies are inline policies that are located in that service. You can't use AWS 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, AWS 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

AWS 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 AWS Organizations. AWS Organizations is a service for grouping and centrally managing multiple AWS 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 AWS account root user. For more information about Organizations and SCPs, see Service control policies in the *AWS Organizations User Guide*.
- **Resource control policies (RCPs)** – RCPs are JSON policies that you can use to set the maximum available permissions for resources in your accounts without updating the IAM policies attached to each resource that you own. The RCP limits permissions for resources in member accounts and can impact the effective permissions for identities, including the AWS account root user, regardless of whether they belong to your organization. For more information about Organizations and RCPs, including a list of AWS services that support RCPs, see Resource control policies (RCPs) in the *AWS 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 AWS determines whether to allow a request when multiple policy types are involved, see Policy evaluation logic in the *IAM User Guide*.

How AWS Storage Gateway works with IAM

Before you use IAM to manage access to AWS SGW, learn what IAM features are available to use with AWS SGW.

IAM features you can use with AWS Storage Gateway

IAM feature	AWS SGW support
Identity-based policies	Yes
Resource-based policies	No
Policy actions	Yes
Policy resources	Yes
Policy condition keys (service-specific)	Yes
ACLs	No
ABAC (tags in policies)	Partial
Temporary credentials	Yes
Forward access sessions (FAS)	Yes
Service roles	Yes
Service-linked roles	Yes

To get a high-level view of how AWS SGW and other AWS services work with most IAM features, see AWS services that work with IAM in the *IAM User Guide*.

Identity-based policies for AWS SGW

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 Define custom IAM permissions with customer managed 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 AWS SGW

To view examples of AWS SGW identity-based policies, see Identity-based policy examples for AWS Storage Gateway.

Resource-based policies within AWS SGW

Supports resource-based policies: No

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 AWS services.

To enable cross-account access, you can specify an entire account or IAM entities in another account as the principal in a resource-based policy. Adding a cross-account principal to a resource-based policy is only half of establishing the trust relationship. When the principal and the resource

are in different AWS accounts, an IAM administrator in the trusted account must also grant the principal entity (user or role) permission to access the resource. They grant permission by attaching an identity-based policy to the entity. However, if a resource-based policy grants access to a principal in the same account, no additional identity-based policy is required. For more information, see Cross account resource access in IAM in the *IAM User Guide*.

Policy actions for AWS SGW

Supports policy actions: Yes

Administrators can use AWS 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 AWS 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 AWS SGW actions, see Actions Defined by AWS Storage Gateway in the *Service Authorization Reference*.

Policy actions in AWS SGW use the following prefix before the action:

```
sgw
```

To specify multiple actions in a single statement, separate them with commas.

```
"Action": [
      "sgw:action1",
      "sgw:action2"
         ]
```

To view examples of AWS SGW identity-based policies, see Identity-based policy examples for AWS Storage Gateway.

Policy resources for AWS SGW

Supports policy resources: Yes

Administrators can use AWS 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 AWS SGW resource types and their ARNs, see Resources Defined by AWS Storage Gateway in the *Service Authorization Reference*. To learn with which actions you can specify the ARN of each resource, see Actions Defined by AWS Storage Gateway .

To view examples of AWS SGW identity-based policies, see Identity-based policy examples for AWS Storage Gateway.

Policy condition keys for AWS SGW

Supports service-specific policy condition keys: Yes

Administrators can use AWS 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, AWS evaluates them using a logical `AND` operation. If you specify multiple values for a single condition key, AWS 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*.

AWS supports global condition keys and service-specific condition keys. To see all AWS global condition keys, see AWS global condition context keys in the *IAM User Guide*.

To see a list of AWS SGW condition keys, see Condition Keys for AWS Storage Gateway in the *Service Authorization Reference*. To learn with which actions and resources you can use a condition key, see Actions Defined by AWS Storage Gateway .

To view examples of AWS SGW identity-based policies, see Identity-based policy examples for AWS Storage Gateway.

ACLs in AWS SGW

Supports ACLs: No

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.

ABAC with AWS SGW

Supports ABAC (tags in policies): Partial

Attribute-based access control (ABAC) is an authorization strategy that defines permissions based on attributes. In AWS, these attributes are called *tags*. You can attach tags to IAM entities (users or roles) and to many AWS 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 Define permissions with ABAC authorization 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*.

Using temporary credentials with AWS SGW

Supports temporary credentials: Yes

Some AWS services don't work when you sign in using temporary credentials. For additional information, including which AWS services work with temporary credentials, see AWS services that work with IAM in the *IAM User Guide*.

You are using temporary credentials if you sign in to the AWS Management Console using any method except a user name and password. For example, when you access AWS 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 Switch from a user to an IAM role (console) in the *IAM User Guide*.

You can manually create temporary credentials using the AWS CLI or AWS API. You can then use those temporary credentials to access AWS. AWS 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 AWS SGW

Supports forward access sessions (FAS): Yes

When you use an IAM user or role to perform actions in AWS, 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 AWS service, combined with the requesting AWS service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other AWS 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 AWS SGW

Supports service roles: Yes

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 Create a role to delegate permissions to an AWS service in the *IAM User Guide*.

> **Warning**
>
> Changing the permissions for a service role might break AWS SGW functionality. Edit service roles only when AWS SGW provides guidance to do so.

Service-linked roles for AWS SGW

Supports service-linked roles: Yes

A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your AWS 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 service-linked roles, see AWS services that work with IAM. Find a service in the table that includes a `Yes` in the **Service-linked role** column. Choose the **Yes** link to view the service-linked role documentation for that service.

Identity-based policy examples for AWS Storage Gateway

By default, users and roles don't have permission to create or modify AWS SGW resources. They also can't perform tasks by using the AWS Management Console, AWS Command Line Interface (AWS CLI), or AWS 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 Create IAM policies (console) in the *IAM User Guide*.

For details about actions and resource types defined by AWS SGW, including the format of the ARNs for each of the resource types, see Actions, Resources, and Condition Keys for AWS Storage Gateway in the *Service Authorization Reference*.

Topics

- Policy best practices
- Using the AWS SGW console
- Allow users to view their own permissions

Policy best practices

Identity-based policies determine whether someone can create, access, or delete AWS SGW resources in your account. These actions can incur costs for your AWS account. When you create or edit identity-based policies, follow these guidelines and recommendations:

- **Get started with AWS managed policies and move toward least-privilege permissions** – To get started granting permissions to your users and workloads, use the *AWS managed policies* that grant permissions for many common use cases. They are available in your AWS account. We recommend that you reduce permissions further by defining AWS customer managed policies that are specific to your use cases. For more information, see AWS managed policies or AWS 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 AWS service, such as AWS 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 Validate policies with IAM Access Analyzer 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 AWS 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 Secure API access with MFA 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 AWS SGW console

To access the AWS Storage Gateway console, you must have a minimum set of permissions. These permissions must allow you to list and view details about the AWS SGW resources in your AWS 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 AWS CLI or the AWS 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 AWS SGW console, also attach the AWS SGW *ConsoleAccess* or *ReadOnly* AWS managed policy to the entities. For more information, see Adding permissions to a user in the *IAM User Guide*.

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 AWS CLI or AWS API.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "ViewOwnUserInfo",
            "Effect": "Allow",
            "Action": [
                "iam:GetUserPolicy",
                "iam:ListGroupsForUser",
                "iam:ListAttachedUserPolicies",
                "iam:ListUserPolicies",
                "iam:GetUser"
            ],
            "Resource": ["arn:aws: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 AWS Storage Gateway identity and access

Use the following information to help you diagnose and fix common issues that you might encounter when working with AWS SGW and IAM.

Topics

- I am not authorized to perform an action in AWS SGW
- I am not authorized to perform iam:PassRole
- I want to allow people outside of my AWS account to access my AWS SGW resources

I am not authorized to perform an action in AWS SGW

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 `sgw:`*`GetWidget`* permissions.

```
User: arn:aws:iam::123456789012:user/mateojackson is not authorized to perform:
 sgw: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 `sgw:GetWidget` action.

If you need help, contact your AWS 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 AWS SGW.

Some AWS 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 AWS SGW. 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: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 AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

Important

Storage Gateway can assume existing service roles that are passed using the `iam:PassRole` policy action, but it does not support IAM policies that use the `iam:PassedToService` context key to limit the action to specific services.
For more information, see the following topics in the *AWS Identity and Access Management User Guide*:

- IAM: Pass an IAM role to a specific AWS service
- Granting a user permissions to pass a role to an AWS service
- Available keys for IAM

I want to allow people outside of my AWS account to access my AWS SGW 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 AWS SGW supports these features, see How AWS Storage Gateway works with IAM.
- To learn how to provide access to your resources across AWS accounts that you own, see Providing access to an IAM user in another AWS account that you own in the *IAM User Guide*.
- To learn how to provide access to your resources to third-party AWS accounts, see Providing access to AWS 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 Cross account resource access in IAM in the *IAM User Guide*.

Using tags to control access to your gateway and resources

To control access to gateway resources and actions, you can use AWS Identity and Access Management (IAM) policies based on tags. You can provide the control in two ways:

1. Control access to gateway resources based on the tags on those resources.
2. Control what tags can be passed in an IAM request condition.

For information about how to use tags to control access, see Controlling Access Using Tags.

Controlling Access Based on Tags on a Resource

To control what actions a user or role can perform on a gateway resource, you can use tags on the gateway resource. For example, you might want to allow or deny specific API operations on a file gateway resource based on the key-value pair of the tag on the resource.

The following example allows a user or a role to perform the `ListTagsForResource`, `ListFileShares`, and `DescribeNFSFileShares` actions on all resources. The policy applies only if the tag on the resource has its key set to `allowListAndDescribe` and the value set to `yes`.

JSON

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "storagegateway:ListTagsForResource",
                "storagegateway:ListFileShares",
                "storagegateway:DescribeNFSFileShares"
            ],
            "Resource": "*",
            "Condition": {
                "StringEquals": {
                    "aws:ResourceTag/allowListAndDescribe": "yes"
                }
            }
        },
        {
            "Effect": "Allow",
            "Action": [
                "storagegateway:*"
            ],
            "Resource": "arn:aws:storagegateway:us-east-1:111122223333:*/*"
        }
    ]
}
```

Controlling Access Based on Tags in an IAM Request

To control what an user can do on a gateway resource, you can use conditions in an IAM policy based on tags. For example, you can write a policy that allows or denies an user the ability to perform specific API operations based on the tag they provided when they created the resource.

In the following example, the first statement allows a user to create a gateway only if the key-value pair of the tag they provided when creating the gateway is **`Department`** and **`Finance`**. When using the API operation, you add this tag to the activation request.

The second statement allows the user to create an Network File System (NFS) or Server Message Block (SMB) file share on a gateway only if the key-value pair of the tag on the gateway matches **`Department`**and **`Finance`**. Additionally, the user must add a tag to the file share, and the key-value pair of the tag must be **`Department`** and **`Finance`**. You add tags to a file share when creating the file share. There aren't permissions for the `AddTagsToResource` or `RemoveTagsFromResource` operations, so the user can't perform these operations on the gateway or the file share.

JSON

```
{
   "Version":"2012-10-17",
   "Statement":[
      {
         "Effect":"Allow",
         "Action":[
            "storagegateway:ActivateGateway"
         ],
         "Resource":"*",
         "Condition":{
            "StringEquals":{
               "aws:RequestTag/Department":"Finance"
            }
         }
      },
      {
         "Effect":"Allow",
         "Action":[
            "storagegateway:CreateNFSFileShare",
            "storagegateway:CreateSMBFileShare"
         ],
         "Resource":"*",
         "Condition":{
            "StringEquals":{
               "aws:ResourceTag/Department":"Finance",
               "aws:RequestTag/Department":"Finance"
            }
```

```
            }
        }
    ]
}
```

Compliance validation for AWS Storage Gateway

Third-party auditors assess the security and compliance of AWS Storage Gateway as part of multiple AWS compliance programs. These include SOC, PCI, ISO, FedRAMP, HIPAA, MTCS, C5, K-ISMS, ENS High, OSPAR, and HITRUST CSF.

For a list of AWS services in scope of specific compliance programs, see AWS Services in Scope by Compliance Program. For general information, see AWS Compliance Programs.

You can download third-party audit reports using AWS Artifact. For more information, see Downloading Reports in AWS Artifact.

Your compliance responsibility when using Storage Gateway is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. AWS 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 security- and compliance-focused baseline environments on AWS.
- Architecting for HIPAA Security and Compliance Whitepaper – This whitepaper describes how companies can use AWS to create HIPAA-compliant applications.
- AWS Compliance Resources – This collection of workbooks and guides might apply to your industry and location.
- Evaluating resources with rules in the *AWS Config Developer Guide* – The AWS Config service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- AWS Security Hub – This AWS service provides a comprehensive view of your security state within AWS that helps you check your compliance with security industry standards and best practices.

Resilience in AWS Storage Gateway

The AWS global infrastructure is built around AWS Regions and Availability Zones.

An AWS Region is a physical location around the world where data centers are clustered. Each group of logical data centers is called an Availability Zone (AZ). Each AWS Region consists of a minimum of three isolated and physically separate AZs within a geographic area. Unlike other cloud providers, who often define a region as a single data center, the multiple AZ design of every AWS Region offers distinct advantages. Each AZ has independent power, cooling, and physical security and is connected via redundant, ultra-low-latency networks. If your deployment requires a focus on high availability, you can configure services and resources to in multiple AZs to achieve greater fault-tolerance.

AWS Regions meet the highest levels of infrastructure security, compliance, and data protection. All traffic between AZs is encrypted. The network performance is sufficient to accomplish synchronous replication between AZs. AZs make partitioning services and resources for high availability easy. If your deployment is partitioned across AZs, your resources are better isolated and protected from issues such as power outages, lightning strikes, tornadoes, earthquakes, and more. AZs are physically separated by a meaningful distance from any other AZ, although all are within 100 km (60 miles) of each other.

For more information about AWS Regions and Availability Zones, see AWS Global Infrastructure.

In addition to the AWS global infrastructure, Storage Gateway supports VMware vSphere High Availability (VMware HA) to help protect storage workloads against hardware, hypervisor, or network failures. For more information, see Using VMware vSphere High Availability with Storage Gateway.

Infrastructure security in AWS Storage Gateway

As a managed service, AWS Storage Gateway is protected by the AWS global network security procedures that are described in Security Pillar - AWS Well-Architected Framework.

You use AWS published API calls to access Storage Gateway through the network. Clients must support Transport Layer Security (TLS) 1.2. Clients must also support cipher suites with perfect forward secrecy (PFS) such as Ephemeral Diffie-Hellman (DHE) or Elliptic Curve Ephemeral Diffie-Hellman (ECDHE). 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 AWS Security Token Service (AWS STS) to generate temporary security credentials to sign requests.

> **Note**
>
> You should treat the AWS Storage Gateway appliance as a managed virtual machine, and should not attempt to access or modify its installation in any way. Attempting to install scanning software or update any software packages using methods other than the normal gateway update mechanism, may cause the gateway to malfunction and could impact our ability to support or fix the gateway.
> AWS reviews, analyzes, and remediates CVEs on a regular basis. We incorporate fixes for these issues into Storage Gateway as part of our normal software release cycle. These fixes are typically applied as part of the normal gateway update process during scheduled maintenance windows. For more information about gateway updates, see Managing gateway updates using the AWS Storage Gateway console.

AWS Security Best Practices

AWS provides a number of security features to consider as you develop and implement your own security policies. These best practices are general guidelines and don't represent a complete security solution. Because these practices might not be appropriate or sufficient for your environment, treat them as helpful considerations rather than prescriptions. For more information, see AWS Security Best Practices.

Logging and monitoring in AWS Storage Gateway

Storage Gateway is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in Storage Gateway. CloudTrail captures all API calls for Storage Gateway as events. The calls captured include calls from the Storage Gateway console and code calls to the Storage Gateway API operations. If you create a trail, you can turn on continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for Storage Gateway. 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 Storage Gateway, 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 AWS CloudTrail User Guide.

Storage Gateway information in CloudTrail

CloudTrail is activated on your AWS account when you create the account. When activity occurs in Storage Gateway, that activity is recorded in a CloudTrail event along with other AWS service events in **Event history**. You can view, search, and download recent events in your AWS account. For more information, see Viewing Events with CloudTrail Event History.

For an ongoing record of events in your AWS account, including events for Storage Gateway, create a trail. A *trail* allows 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 AWS Regions. The trail logs events from all Regions in the AWS partition and delivers the log files to the Amazon S3 bucket that you specify. Additionally, you can configure other AWS services to further analyze and act upon the event data collected in CloudTrail logs. For more information, see the following:

- Overview for Creating a Trail
- CloudTrail Supported Services and Integrations
- Configuring Amazon SNS Notifications for CloudTrail
- Receiving CloudTrail Log Files from Multiple Regions and Receiving CloudTrail Log Files from Multiple Accounts

All of the Storage Gateway actions are logged and are documented in the Actions topic. For example, calls to the `ActivateGateway`, `ListGateways`, and `ShutdownGateway` actions 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 AWS 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 AWS service.

For more information, see the CloudTrail userIdentity Element.

Understanding Storage Gateway log file entries

A trail is a configuration that allows 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 action.

```
{ "Records": [{
               "eventVersion": "1.02",
               "userIdentity": {
               "type": "IAMUser",
               "principalId": "AIDAII5AUEPBH2M7JTNVC",
               "arn": "arn:aws:iam::111122223333:user/StorageGateway-team/JohnDoe",
               "accountId": "111122223333",
               "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
                "userName": "JohnDoe"
              },
                 "eventTime": "2014-12-04T16:19:00Z",
                 "eventSource": "storagegateway.amazonaws.com",
                 "eventName": "ActivateGateway",
                 "awsRegion": "us-east-2",
                 "sourceIPAddress": "192.0.2.0",
                 "userAgent": "aws-cli/1.6.2 Python/2.7.6 Linux/2.6.18-164.el5",
                  "requestParameters": {
                                          "gatewayTimezone": "GMT-5:00",
                                          "gatewayName": "cloudtrailgatewayvtl",
                                          "gatewayRegion": "us-east-2",
                                          "activationKey": "EHFBX-1NDD0-P0IVU-PI259-
DHK88",
                                          "gatewayType": "VTL"
                                               },
                                               "responseElements": {
                                                                       "gatewayARN":
 "arn:aws:storagegateway:us-east-2:111122223333:gateway/cloudtrailgatewayvtl"
                                               },
                                               "requestID":
 "54BTFGNQI71987UJD2IHTCT8NF1Q8GLLE1QEU3KPGG6F0KSTAUU0",
                                               "eventID": "635f2ea2-7e42-45f0-
bed1-8b17d7b74265",
                                               "eventType": "AwsApiCall",
```

```
                                                    "apiVersion": "20130630",
                                                    "recipientAccountId": "444455556666"
            }]
}
```

The following example shows a CloudTrail log entry that demonstrates the ListGateways action.

```
{
 "Records": [{
               "eventVersion": "1.02",
               "userIdentity": {
                                "type": "IAMUser",
                                "principalId": "AIDAII5AUEPBH2M7JTNVC",
                                "arn": "arn:aws:iam::111122223333:user/StorageGateway-
team/JohnDoe",
                                "accountId:" 111122223333", " accessKeyId ":"
 AKIAIOSFODNN7EXAMPLE",
                                " userName ":" JohnDoe "
                                },

                                " eventTime ":" 2014 - 12 - 03T19: 41: 53Z ",
                                " eventSource ":" storagegateway.amazonaws.com ",
                                " eventName ":" ListGateways ",
                                " awsRegion ":" us-east-2 ",
                                " sourceIPAddress ":" 192.0.2.0 ",
                                " userAgent ":" aws - cli / 1.6.2 Python / 2.7.6
 Linux / 2.6.18 - 164.el5 ",
                                " requestParameters ":null,
                                " responseElements ":null,
                                "requestID ":"
 6U2N42CU37KAO8BG6V1I23FRSJ1Q8GLLE1QEU3KPGG6F0KSTAUU0 ",
                                " eventID ":" f76e5919 - 9362 - 48ff - a7c4 -
 d203a189ec8d ",
                                " eventType ":" AwsApiCall ",
                                " apiVersion ":" 20130630 ",
                                " recipientAccountId ":" 444455556666"
              }]
}
```

Troubleshooting problems with your Storage Gateway deployment

Following, you can find information about best practices and troubleshooting issues related to gateways, host platforms, file systems, high availability, data recovery, and snapshots. The on-premises gateway troubleshooting information covers gateways deployed on supported virtualization platforms. The troubleshooting information for high availability issues covers gateways running on VMware vSphere High Availability (HA) platform.

Topics

- Troubleshooting: gateway offline issues - Learn how to diagnose problems that can cause your gateway to appear offline in the Storage Gateway console.
- Troubleshooting: Active Directory issues - Learn what to do if you receive error messages such as `NETWORK_ERROR`, `TIMEOUT`, or `ACCESS_DENIED` when trying to join your File Gateway to a Microsoft Active Directory domain.
- Troubleshooting: gateway activation issues - Learn what to do if you receive an internal error message when attempting to activate your Storage Gateway.
- Troubleshooting: on-premises gateway issues - Learn about typical issues that you might encounter working with your on-premises gateways, and how to allow Support to connect to your gateway to assist with troubleshooting.
- Troubleshooting: Microsoft Hyper-V setup issues - Learn about typical issues that you might encounter when deploying Storage Gateway on the Microsoft Hyper-V platform.
- Troubleshooting: Amazon EC2 gateway issues - Find information about typical issues that you might encounter when working with gateways deployed on Amazon EC2.
- Troubleshooting: hardware appliance issues - Learn how to resolve issues that you might encounter with the AWS Storage Gateway Hardware Appliance.
- Troubleshooting: File Gateway issues - Find information that can help you understand the cause of errors and health notifications that appear in your File Gateway's CloudWatch logs.
- Troubleshooting: high availability issues - Learn what to do if you experience issues with gateways that are deployed in a VMware HA environment.

Troubleshooting: gateway offline in the Storage Gateway console

Use the following troubleshooting information to determine what to do if the AWS Storage Gateway console shows that your gateway is offline.

Your gateway might be showing as offline for one or more of the following reasons:

- The gateway can't reach the Storage Gateway service endpoints.
- The gateway shut down unexpectedly.
- A cache disk associated with the gateway has been disconnected or modified, or has failed.

To bring your gateway back online, identify and resolve the issue that caused your gateway to go offline.

Check the associated firewall or proxy

If you configured your gateway to use a proxy, or you placed your gateway behind a firewall, then review the access rules of the proxy or firewall. The proxy or firewall must allow traffic to and from the network ports and service endpoints required by Storage Gateway. For more information, see Network and firewall requirements.

Check for an ongoing SSL or deep-packet inspection of your gateway's traffic

If an SSL or deep-packet inspection is currently being performed on the network traffic between your gateway and AWS, then your gateway might not be able to communicate with the required service endpoints. To bring your gateway back online, you must disable the inspection.

Check the IOWaitPercent metric after a reboot or software update

After a reboot or software update, check to see if the `IOWaitPercent` metric for your File Gateway is 10 or greater. This might cause your gateway to be slow to respond while it rebuilds the index cache to RAM. For more information, see Troubleshooting: Using CloudWatch metrics.

Check for a power outage or hardware failure on the hypervisor host

A power outage or hardware failure on the hypervisor host of your gateway can cause your gateway to shut down unexpectedly and become unreachable. After you restore the power and network connectivity, your gateway will become reachable again.

After your gateway is back online, be sure to take steps to recover your data. For more information, see Best practices: recovering your data.

Check for issues with an associated cache disk

Your gateway can go offline if at least one of the cache disks associated with your gateway was removed, changed, or resized, or if it is corrupted.

If a working cache disk was removed from the hypervisor host:

1. Shut down the gateway.
2. Re-add the disk.

   > **Note**
   >
   > Make sure you add the disk to the same disk node.

3. Restart the gateway.

If a cache disk is corrupted, was replaced, or was resized:

- Follow the **Method 2** procedure described in Replacing your existing S3 File Gateway with a new instance to set up a new gateway and re-download cache disk information from the AWS cloud.

Troubleshooting: issues joining gateway to Active Directory

Use the following troubleshooting information to determine what to do if you receive error messages such as `NETWORK_ERROR`, `TIMEOUT`, or `ACCESS_DENIED` when trying to join your File Gateway to a Microsoft Active Directory domain.

To resolve these errors, perform the following checks and configurations.

Confirm that the gateway can reach the domain controller by running an nping test

To run an nping test:

1. Connect to the gateway local console using your hypervisor management software (VMware, Hyper-V, or KVM) for on-premises gateways, or using ssh for Amazon EC2 gateways.
2. Enter the corresponding numeral to select **Gateway Console**, and then enter h to list all available commands. To test the connectivity between the Storage Gateway virtual machine and the domain, run the following command:

   ```
   nping -d corp.domain.com -p 389 -c 1 -t tcp
   ```

   > **Note**
   >
   > Replace `corp.domain.com` with your Active Directory domain DNS name and replace 389 with the LDAP port for your environment.
   > Verify that you have opened the required ports within your firewall.

The following is an example of a successful nping test where the gateway was able to reach the domain controller:

```
nping -d corp.domain.com -p 389 -c 1 -t tcp

Starting Nping 0.6.40 ( http://nmap.org/nping ) at 2022-06-30 16:24 UTC
SENT (0.0553s) TCP 10.10.10.21:9783 > 10.10.10.10:389 S ttl=64 id=730 iplen=40
 seq=2597195024 win=1480
RCVD (0.0556s) TCP 10.10.10.10:389 > 10.10.10.21:9783 SA ttl=128 id=22332 iplen=44
 seq=4170716243 win=8192 <mss 8961>

Max rtt: 0.310ms | Min rtt: 0.310ms | Avg rtt: 0.310ms
Raw packets sent: 1 (40B) | Rcvd: 1 (44B) | Lost: 0 (0.00%)
Nping done: 1 IP address pinged in 1.09 seconds<br>
```

The following is an example of an nping test where there is no connectivity to or response from the `corp.domain.com` destination:

```
nping -d corp.domain.com -p 389 -c 1 -t tcp
```

```
Starting Nping 0.6.40 ( http://nmap.org/nping ) at 2022-06-30 16:26 UTC
SENT (0.0421s) TCP 10.10.10.21:47196 > 10.10.10.10:389  S ttl=64 id=30318 iplen=40
 seq=1762671338 win=1480

Max rtt: N/A | Min rtt: N/A | Avg rtt: N/A
Raw packets sent: 1 (40B) | Rcvd: 0 (0B) | Lost: 1 (100.00%)
Nping done: 1 IP address pinged in 1.07 seconds
```

Check the DHCP options set for the VPC of your Amazon EC2 gateway instance

If the File Gateway is running on an Amazon EC2 instance, then you must make sure a DHCP options set is properly configured and attached to the Amazon Virtual Private Cloud (VPC) that contains the gateway instance. For more information, see DHCP option sets in Amazon VPC.

Confirm that the gateway can resolve the domain by running a dig query

If the domain isn't resolvable by the gateway, then the gateway can't join the domain.

To run a dig query:

1. Connect to the gateway local console using your hypervisor management software (VMware, Hyper-V, or KVM) for on-premises gateways, or using ssh for Amazon EC2 gateways.
2. Enter the corresponding numeral to select **Gateway Console**, and then enter h to list all available commands. To test whether the gateway can resolve the domain, run the following command:

   ```
   dig -d corp.domain.com
   ```

   > **Note**
   >
   > Replace `corp.domain.com` with your Active Directory domain DNS name.

The following is an example of a successful response:

```
; <<>> DiG 9.11.4-P2-RedHat-9.11.4-26.P2.amzn2.5.2 <<>> corp.domain.com
```

```
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 24817
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 2, AUTHORITY: 0, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 4000
;; QUESTION SECTION:
;corp.domain.com.        IN    A

;; ANSWER SECTION:
corp.domain.com.    600    IN    A    10.10.10.10
corp.domain.com.    600    IN    A    10.10.20.10

;; Query time: 0 msec
;; SERVER: 10.10.20.228#53(10.10.20.228)
;; WHEN: Thu Jun 30 16:36:32 UTC 2022
;; MSG SIZE  rcvd: 78
```

Check the domain controller settings and roles

Verify that the domain controller isn't set to read-only, and that the domain controller has enough roles for computers to join. To test this, try joining other servers from the same VPC subnet as the gateway VM to the domain.

Check that the gateway is joined to the nearest domain controller

As a best practice, we recommend joining your gateway to a domain controller that is geographically close to the gateway appliance. If the gateway appliance can't communicate with the domain controller within 20 seconds due to network latency, then the domain joining process can time out. For example, the process might time out if the gateway appliance is in the US East (N. Virginia) AWS Region and the domain controller is in the Asia Pacific (Singapore) AWS Region.

> **Note**
>
> To increase the default timeout value of 20 seconds, you can run the join-domain command in the AWS Command Line Interface (AWS CLI) and include the `--timeout-in-seconds` option to increase the time. You can also use the JoinDomain API call and include the `TimeoutInSeconds` parameter to increase the time. The maximum timeout value is 3,600 seconds.

If you receive errors when running AWS CLI commands, make sure that you're using the most recent AWS CLI version.

Confirm that Active Directory creates new computer objects in the default organizational unit (OU)

Make sure Microsoft Active Directory does not have any Group Policy Objects that create new computer objects in any location other than the default OU. Before you can join your gateway to the Active Directory domain, a new computer object must exist in the default OU. Some Active Directory environments are customized to have different OUs for newly created objects. To guarantee that a new computer object for the gateway VM exists in the default OU, try creating the computer object manually on your domain controller before you join the gateway to the domain. You can also run the join-domain command using the AWS CLI. Then, specify the option for `--organizational-unit`.

Note

The process of creating the computer object is called pre-staging.

Check your domain controller event logs

If you can't join the gateway to the domain after trying all other checks and configurations described in the previous sections, we recommend examining your domain controller event logs. Check for any errors in the event viewer of the domain controller. Verify that the gateway queries have reached the domain controller.

Troubleshooting: internal error during gateway activation

Storage Gateway activation requests traverse two network paths. Incoming activation requests sent by a client connect to the gateway's virtual machine (VM) or Amazon Elastic Compute Cloud (Amazon EC2) instance over port 80. If the gateway successfully receives the activation request, then the gateway communicates with the Storage Gateway endpoints to receive an activation key. If the gateway can't reach the Storage Gateway endpoints, then the gateway responds to the client with an internal error message.

Use the following troubleshooting information to determine what to do if you receive an internal error message when attempting to activate your AWS Storage Gateway.

Note

- Make sure you deploy new gateways using the latest virtual machine image file or Amazon Machine Image (AMI) version. You will receive an internal error if you attempt to activate a gateway that uses an outdated AMI.
- Make sure that you select the correct gateway type that you intend to deploy before you download the AMI. The .ova files and AMIs for each gateway type are different, and they are not interchangeable.

Resolve errors when activating your gateway using a public endpoint

To resolve activation errors when activating your gateway using a public endpoint, perform the following checks and configurations.

Check the required ports

For gateways deployed on-premises, check that the ports are open on your local firewall. For gateways deployed on an Amazon EC2 instance, check that the ports are open on the instance's security group. To confirm that the ports are open, run a telnet command on the public endpoint from a server. This server must be in the same subnet as the gateway. For example, the following telnet commands test the connection to port 443:

```
telnet d4kdq0yaxexbo.cloudfront.net 443
telnet storagegateway.region.amazonaws.com 443
telnet dp-1.storagegateway.region.amazonaws.com 443
telnet proxy-app.storagegateway.region.amazonaws.com 443
telnet client-cp.storagegateway.region.amazonaws.com 443
telnet anon-cp.storagegateway.region.amazonaws.com 443
```

To confirm that the gateway itself can reach the endpoint, access the gateway's local VM console (for gateways deployed on-premises). Or, you can SSH to the gateway's instance (for gateways deployed on Amazon EC2). Then, run a network connectivity test. Confirm that the test returns `[PASSED]`. For more information, see Testing your gateway's network connectivity.

> **Note**
>
> The default login user name for the gateway console is `admin`, and the default password is `password`.

Make sure firewall security does not modify packets sent from the gateway to the public endpoints

SSL inspections, deep packet inspections, or other forms of firewall security can interfere with packets sent from the gateway. The SSL handshake fails if the SSL certificate is modified from what the activation endpoint expects. To confirm that there's no SSL inspection in progress, run an OpenSSL command on the main activation endpoint (`anon-cp.storagegateway.region.amazonaws.com`) on port 443. You must run this command from a machine that's in the same subnet as the gateway:

```
$ openssl s_client -connect  anon-cp.storagegateway.region.amazonaws.com:443 -
servername anon-cp.storagegateway.region.amazonaws.com
```

> **Note**
>
> Replace *region* with your AWS Region.

If there's no SSL inspection in progress, then the command returns a response similar to the following:

```
$ openssl s_client -connect anon-cp.storagegateway.us-east-2.amazonaws.com:443 -
servername anon-cp.storagegateway.us-east-2.amazonaws.com
CONNECTED(00000003)
depth=2 C = US, O = Amazon, CN = Amazon Root CA 1
verify return:1
depth=1 C = US, O = Amazon, OU = Server CA 1B, CN = Amazon
verify return:1
depth=0 CN = anon-cp.storagegateway.us-east-2.amazonaws.com
verify return:1
---
Certificate chain
 0 s:/CN=anon-cp.storagegateway.us-east-2.amazonaws.com
```

```
   i:/C=US/O=Amazon/OU=Server CA 1B/CN=Amazon
 1 s:/C=US/O=Amazon/OU=Server CA 1B/CN=Amazon
   i:/C=US/O=Amazon/CN=Amazon Root CA 1
 2 s:/C=US/O=Amazon/CN=Amazon Root CA 1
   i:/C=US/ST=Arizona/L=Scottsdale/O=Starfield Technologies, Inc./CN=Starfield Services
 Root Certificate Authority - G2
 3 s:/C=US/ST=Arizona/L=Scottsdale/O=Starfield Technologies, Inc./CN=Starfield Services
 Root Certificate Authority - G2
   i:/C=US/O=Starfield Technologies, Inc./OU=Starfield Class 2 Certification Authority
---
```

If there is an ongoing SSL inspection, then the response shows an altered certificate chain, similar to the following:

```
$ openssl s_client -connect  anon-cp.storagegateway.ap-southeast-1.amazonaws.com:443 -
servername anon-cp.storagegateway.ap-southeast-1.amazonaws.com
CONNECTED(00000003)
depth=0 DC = com, DC = amazonaws, OU = AWS, CN = anon-cp.storagegateway.ap-
southeast-1.amazonaws.com
verify error:num=20:unable to get local issuer certificate
verify return:1
depth=0 DC = com, DC = amazonaws, OU = AWS, CN = anon-cp.storagegateway.ap-
southeast-1.amazonaws.com
verify error:num=21:unable to verify the first certificate
verify return:1
---
Certificate chain
 0 s:/DC=com/DC=amazonaws/OU=AWS/CN=anon-cp.storagegateway.ap-southeast-1.amazonaws.com
   i:/C=IN/O=Company/CN=Admin/ST=KA/L=New town/OU=SGW/emailAddress=admin@company.com
---
```

The activation endpoint accepts SSL handshakes only if it recognizes the SSL certificate. This means that the gateway's outbound traffic to the endpoints must be exempt from inspections performed by firewalls in your network. These inspections might be an SSL inspection or a deep packet inspection.

Check gateway time synchronization

Excessive time skews can cause SSL handshake errors. For on-premises gateways, you can use the gateway's local VM console to check your gateway's time synchronization. The time skew should be no larger than 60 seconds. For more information, see Synchronizing Your Gateway VM Time.

The **System Time Management** option isn't available on gateways that are hosted on Amazon EC2 instances. To make sure Amazon EC2 gateways can properly synchronize time, confirm that the Amazon EC2 instance can connect to the following NTP server pool list over ports UDP and TCP 123:

- 0.amazon.pool.ntp.org
- 1.amazon.pool.ntp.org
- 2.amazon.pool.ntp.org
- 3.amazon.pool.ntp.org

Resolve errors when activating your gateway using an Amazon VPC endpoint

To resolve activation errors when activating your gateway using an Amazon Virtual Private Cloud (Amazon VPC) endpoint, perform the following checks and configurations.

Check the required ports

Make sure the required ports within your local firewall (for gateways deployed on-premises) or security group (for gateways deployed in Amazon EC2) are open. The ports required for connecting a gateway to a Storage Gateway VPC endpoint differ from those required when connecting a gateway to public endpoints. The following ports are required for connecting to a Storage Gateway VPC endpoint:

- TCP 443
- TCP 1026
- TCP 1027
- TCP 1028
- TCP 1031
- TCP 2222

For more information, see Creating a VPC endpoint for Storage Gateway.

Additionally, check the security group that's attached to your Storage Gateway VPC endpoint. The default security group attached to the endpoint might not allow the required ports. Create a new

security group that allows traffic from your gateway's IP address range over the required ports. Then, attach that security group to the VPC endpoint.

> **Note**
>
> Use the Amazon VPC console to verify the security group that's attached to the VPC endpoint. View your Storage Gateway VPC endpoint from the console, and then choose the **Security Groups** tab.

To confirm that the required ports are open, you can run telnet commands on the Storage Gateway VPC Endpoint. You must run these commands from a server that's in the same subnet as the gateway. You can run the tests on the first DNS name that doesn't specify an Availability Zone. For example, the following telnet commands test the required port connections using the DNS name vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com:

```
telnet vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com 443
telnet vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com 1026
telnet vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com 1027
telnet vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com 1028
telnet vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com 1031
telnet vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com 2222
```

Make sure firewall security does not modify packets sent from the gateway to your Storage Gateway Amazon VPC endpoint

SSL inspections, deep packet inspections, or other forms of firewall security can interfere with packets sent from the gateway. The SSL handshake fails if the SSL certificate is modified from what the activation endpoint expects. To confirm that there's no SSL inspection in progress, run an OpenSSL command on your Storage Gateway VPC endpoint. You must run this command from a machine that's in the same subnet as the gateway. Run the command for each required port:

```
$ openssl s_client -connect vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-
east-1.vpce.amazonaws.com:443 -servername
 vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com

$ openssl s_client -connect vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-
east-1.vpce.amazonaws.com:1026 -servername
 vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com
```

```
$ openssl s_client -connect vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-
east-1.vpce.amazonaws.com:1027 -servername
 vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com

$ openssl s_client -connect vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-
east-1.vpce.amazonaws.com:1028 -servername
 vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com

$ openssl s_client -connect vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-
east-1.vpce.amazonaws.com:1031 -servername
 vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com

$ openssl s_client -connect vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-
east-1.vpce.amazonaws.com:2222 -servername
 vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com
```

If there's no SSL inspection in progress, then the command returns a response similar to the following:

```
openssl s_client -connect vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-
east-1.vpce.amazonaws.com:1027 -servername
 vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com
CONNECTED(00000005)
depth=2 C = US, O = Amazon, CN = Amazon Root CA 1
verify return:1
depth=1 C = US, O = Amazon, OU = Server CA 1B, CN = Amazon
verify return:1
depth=0 CN = anon-cp.storagegateway.us-east-1.amazonaws.com
verify return:1
---
Certificate chain
 0 s:CN = anon-cp.storagegateway.us-east-1.amazonaws.com
   i:C = US, O = Amazon, OU = Server CA 1B, CN = Amazon
 1 s:C = US, O = Amazon, OU = Server CA 1B, CN = Amazon
   i:C = US, O = Amazon, CN = Amazon Root CA 1
 2 s:C = US, O = Amazon, CN = Amazon Root CA 1
   i:C = US, ST = Arizona, L = Scottsdale, O = "Starfield Technologies, Inc.", CN =
 Starfield Services Root Certificate Authority - G2
 3 s:C = US, ST = Arizona, L = Scottsdale, O = "Starfield Technologies, Inc.", CN =
 Starfield Services Root Certificate Authority - G2
   i:C = US, O = "Starfield Technologies, Inc.", OU = Starfield Class 2 Certification
 Authority
---
```

If there is an ongoing SSL inspection, then the response shows an altered certificate chain, similar to the following:

```
openssl s_client -connect vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-
east-1.vpce.amazonaws.com:1027 -servername
 vpce-1234567e1c24a1fe9-62qntt8k.storagegateway.us-east-1.vpce.amazonaws.com
CONNECTED(00000005)
depth=2 C = US, O = Amazon, CN = Amazon Root CA 1
verify return:1
depth=1 C = US, O = Amazon, OU = Server CA 1B, CN = Amazon
verify return:1
depth=0 DC = com, DC = amazonaws, OU = AWS, CN = anon-cp.storagegateway.us-
east-1.amazonaws.com
verify error:num=21:unable to verify the first certificate
verify return:1
---
Certificate chain
 0 s:/DC=com/DC=amazonaws/OU=AWS/CN=anon-cp.storagegateway.us-east-1.amazonaws.com
   i:/C=IN/O=Company/CN=Admin/ST=KA/L=New town/OU=SGW/emailAddress=admin@company.com
---
```

The activation endpoint accepts SSL handshakes only if it recognizes the SSL certificate. This means that the gateway's outbound traffic to your VPC endpoint over required ports is exempt from inspections performed by your network firewalls. These inspections might be SSL inspections or deep packet inspections.

Check gateway time synchronization

Excessive time skews can cause SSL handshake errors. For on-premises gateways, you can use the gateway's local VM console to check your gateway's time synchronization. The time skew should be no larger than 60 seconds. For more information, see Synchronizing Your Gateway VM Time.

The **System Time Management** option isn't available on gateways that are hosted on Amazon EC2 instances. To make sure Amazon EC2 gateways can properly synchronize time, confirm that the Amazon EC2 instance can connect to the following NTP server pool list over ports UDP and TCP 123:

- 0.amazon.pool.ntp.org
- 1.amazon.pool.ntp.org
- 2.amazon.pool.ntp.org

- 3.amazon.pool.ntp.org

Check for an HTTP proxy and confirm associated security group settings

Before activation, check if you have an HTTP proxy on Amazon EC2 configured on the on-premises gateway VM as a Squid proxy on port 3128. In this case, confirm the following:

- The security group attached to the HTTP proxy on Amazon EC2 must have an inbound rule. This inbound rule must allow Squid proxy traffic on port 3128 from the gateway VM's IP address.
- The security group attached to the Amazon EC2 VPC endpoint must have inbound rules. These inbound rules must allow traffic on ports 1026-1028, 1031, 2222, and 443 from the IP address of the HTTP proxy on Amazon EC2.

Resolve errors when activating your gateway using a public endpoint and there is a Storage Gateway VPC endpoint in the same VPC

To resolve errors when activating your gateway using a public endpoint when there is a Amazon Virtual Private Cloud (Amazon VPC) enpoint in the same VPC, perform the following checks and configurations.

Confirm that the Enable Private DNS Name setting isn't enabled on your Storage Gateway VPC endpoint

If **Enable Private DNS Name** is enabled, you can't activate any gateways from that VPC to the public endpoint.

To disable the private DNS name option:

1. Open the Amazon VPC console.
2. In the navigation pane, choose **Endpoints**.
3. Choose your Storage Gateway VPC endpoint.
4. Choose **Actions**.
5. Choose **Manage Private DNS Names**.
6. For **Enable Private DNS Name**, clear **Enable for this Endpoint**.
7. Choose **Modify Private DNS Names** to save the setting.

Troubleshooting: on-premises gateway issues

You can find information following about typical issues that you might encounter working with your on-premises gateways, and how to allow Support to connect to your gateway to assist with troubleshooting.

The following table lists typical issues that you might encounter working with your on-premises gateways.

Issue	Action to Take
You cannot find the IP address of your gateway.	Use the hypervisor client to connect to your host to find the gateway IP address. • For VMware ESXi, the VM's IP address can be found in the vSphere client on the **Summary** tab. • For Microsoft Hyper-V, the VM's IP address can be found by logging into the local console. If you are still having trouble finding the gateway IP address: • Check that the VM is turned on. Only when the VM is turned on does an IP address get assigned to your gateway. • Wait for the VM to finish startup. If you just turned on your VM, then it might take several minutes for the gateway to finish its boot sequence.
You're having network or firewall problems.	• Allow the appropriate ports for your gateway. • If you use a firewall or router to filter or limit network traffic, you must configure your firewall and router to allow these service endpoints for outbound communication to AWS. For more information about network and firewall requirements, see Network and firewall requirements.
Your gateway's activatio n fails when you click the **Proceed to Activation**	• Check that the gateway VM can be accessed by pinging the VM from your client.

Issue	Action to Take
button in the Storage Gateway Management Console.	• Check that your VM has network connectivity to the internet. Otherwise, you'll need to configure a SOCKS proxy. For more information on doing so, see Testing your gateway's network connectivity. • Check that the host has the correct time, that the host is configured to synchronize its time automatically to a Network Time Protocol (NTP) server, and that the gateway VM has the correct time. For information about synchronizing the time of hypervisor hosts and VMs, see Configuring a Network Time Protocol (NTP) server for your gateway. • After performing these steps, you can retry the gateway deployment using the Storage Gateway console and the **Setup and Activate Gateway** wizard. • Check that your VM has at least 16 GB of RAM. Gateway allocation fails if there is less than 16 GB of RAM. For more information, see File Gateway setup requirements.
You need to improve bandwidth between your gateway and AWS.	You can improve the bandwidth from your gateway to AWS by setting up your internet connection to AWS on a network adapter (NIC) separate from that connecting your applications and the gateway VM. Taking this approach is useful if you have a high-bandwidth connection to AWS and you want to avoid bandwidth contention, especially during a snapshot restore. For high-throughput workload needs, you can use AWS Direct Connect to establish a dedicated network connection between your on-premises gateway and AWS. To measure the bandwidth of the connection from your gateway to AWS, use the `CloudBytesDownloaded` and `CloudBytesUploaded` metrics of the gateway. For more on this subject, see Performance and optimization. Improving your internet connectivity helps to ensure that your upload buffer does not fill up.

Issue	Action to Take

Throughput to or from your gateway drops to zero.	• On the **Gateway** tab of the Storage Gateway console, verify that the IP addresses for your gateway VM are the same that you see using your hypervisor client software (that is, the VMware vSphere client or Microsoft Hyper-V Manager). If you find a mismatch, restart your gateway from the Storage Gateway console, as shown in Shutting down your gateway VM. After the restart, the addresses in the **IP Addresses** list in the Storage Gateway console's **Gateway** tab should match the IP addresses for your gateway, which you determine from the hypervisor client.
	  • For VMware ESXi, the VM's IP address can be found in the vSphere client on the **Summary** tab.
	  • For Microsoft Hyper-V, the VM's IP address can be found by logging into the local console.
	• Check your gateway's connectivity to AWS as described in Testing your gateway's network connectivity.
	• Check your gateway's network adapter configuration in your hypervisor management client and ensure that all the interfaces you intend to use for the gateway are activated.
	• Check your gateway's network adapter configuration in the gateway local console. For instructions, see Configuring your gateway network settings.
	
	You can view the throughput to and from your gateway from the Amazon CloudWatch console. For more information about measuring throughput to and from your gateway to AWS, see Performance and optimization.

You are having trouble importing (deploying) Storage Gateway on Microsoft Hyper-V.	See Troubleshooting: Microsoft Hyper-V setup, which discusses some of the common issues of deploying a gateway on Microsoft Hyper-V.

Issue	Action to Take
You receive a message that says: "The data that has been written to the volume in your gateway isn't securely stored at AWS".	You receive this message if your gateway VM was created from a clone or snapshot of another gateway VM. If this isn't the case, contact Support.

Turning on Support access to help troubleshoot your gateway hosted on-premises

Storage Gateway provides a local console you can use to perform several maintenance tasks, including allowing Support to access your gateway to assist you with troubleshooting gateway issues. By default, Support access to your gateway is turned off. You turn on this access through the host's local console. To give Support access to your gateway, you first log in to the local console for the host, navigate to the Storage Gateway's console, and then connect to the support server.

To turn on Support access to your gateway

1. Log in to your host's local console.
   - VMware ESXi – for more information, see Accessing the Gateway Local Console with VMware ESXi.
   - Microsoft Hyper-V – for more information, see Access the Gateway Local Console with Microsoft Hyper-V.
2. At the prompt, enter the corresponding numeral to select **Gateway Console**.
3. Enter **h** to open the list of available commands.
4. Do one of the following:
   - If your gateway is using a public endpoint, in the **AVAILABLE COMMANDS** window, enter **`open-support-channel`** to connect to customer support for Storage Gateway. Allow TCP port 22 so you can open a support channel to AWS. When you connect to customer support, Storage Gateway assigns you a support number. Make a note of your support number.
   - If your gateway is using a VPC endpoint, in the **AVAILABLE COMMANDS** window, enter **`open-support-channel`**. If your gateway is not activated, provide the VPC endpoint or IP

address to connect to customer support for Storage Gateway. Allow TCP port 22 so you can open a support channel to AWS. When you connect to customer support, Storage Gateway assigns you a support number. Make a note of your support number.

> **Note**
>
> The channel number is not a Transmission Control Protocol/User Datagram Protocol (TCP/UDP) port number. Instead, the gateway makes a Secure Shell (SSH) (TCP 22) connection to Storage Gateway servers and provides the support channel for the connection.

5. After the support channel is established, provide your support service number to Support so Support can provide troubleshooting assistance.
6. When the support session is completed, enter **q** to end it. Don't close the session until Amazon Web Services Support notifies you that the support session is complete.
7. Enter **`exit`** to log out of the Storage Gateway console.
8. Follow the prompts to exit the local console.

Troubleshooting: Microsoft Hyper-V setup

The following table lists typical issues that you might encounter when deploying Storage Gateway on the Microsoft Hyper-V platform.

Issue	Action to Take
You try to import a gateway and receive the following error message: "A server error occurred while attempting to import the virtual machine. Import failed. Unable to find virtual machine import files under location [...]. You	This error can occur for the following reasons: • If you are not pointing to the root of the unzipped gateway source files. The last part of the location you specify in the **Import Virtual Machine** dialog box should be `AWS-Storage-Gateway` . For example: `C:\prod-gateway\unzippedSourceVM\AWS-Storage-Gateway\` . • If you have already deployed a gateway and you did not select the **Copy the virtual machine** option and check the **Duplicate**

Issue	Action to Take
can import a virtual machine only if you used Hyper-V to create and export it."	**all files** option in the **Import Virtual Machine** dialog box, then the VM was created in the location where you have the unzipped gateway files and you cannot import from this location again. To fix this problem, get a fresh copy of the unzipped gateway source files and copy to a new location. Use the new location as the source of the import. If you plan on creating multiple gateways from one unzipped source files location, you must select **Copy the virtual machine** and check the **Duplicate all files** box in the **Import Virtual Machine** dialog box.
You try to import a gateway and receive the following error message: "A server error occurred while attempting to import the virtual machine. Import failed. Import task failed to copy file from [...]: The file exists. (0x80070050)"	If you have already deployed a gateway and you try to reuse the default folders that store the virtual hard disk files and virtual machine configuration files, then this error will occur. To fix this problem, specify new locations under **Server** in the panel on the left side of the **Hyper-V Settings** dialog box.

Issue	Action to Take
You try to import a gateway and receive the following error message: "A server error occurred while attempting to import the virtual machine. Import failed. Import failed because the virtual machine must have a new identifier. Select a new identifier and try the import again."	When you import the gateway make sure you select **Copy the virtual machine** and check the **Duplicate all files** box in the **Import Virtual Machine** dialog box to create a new unique ID for the VM.
You try to start a gateway VM and receive the following error message: "An error occurred while attempting to start the selected virtual machine(s). The child partition processor setting is incompatible with parent partition. 'AWS-Storage-Gateway' could not initialize. (Virtual machine ID [...])"	This error is likely caused by a CPU discrepancy between the required CPUs for the gateway and the available CPUs on the host. Ensure that the VM CPU count is supported by the underlying hypervisor. For more information about the requirements for Storage Gateway, see File Gateway setup requirements.

Issue	Action to Take
You try to start a gateway VM and receive the following error message: "An error occurred while attempting to start the selected virtual machine(s). 'AWS-Storage-Gateway' could not initialize. (Virtual machine ID [...]) Failed to create partition: Insufficient system resources exist to complete the requested service. (0x800705AA)"	This error is likely caused by a RAM discrepancy between the required RAM for the gateway and the available RAM on the host. For more information about the requirements for Storage Gateway, see File Gateway setup requirements.
Your snapshots and gateway software updates are occurring at slightly different times than expected.	The gateway VM's clock might be offset from the actual time, known as clock drift. Check and correct the VM's time using local gateway console's time synchronization option. For more information, see Configuring a Network Time Protocol (NTP) server for your gateway.
You need to put the unzipped Microsoft Hyper-V Storage Gateway files on the host file system.	Access the host as you do a typical Microsoft Windows server. For example, if the hypervisor host is name `hyperv-server` , then you can use the following UNC path `\\hyperv-server\c$` , which assumes that the name `hyperv-server` can be resolved or is defined in your local hosts file.
You are prompted for credentials when connecting to hypervisor.	Add your user credentials as a local administrator for the hypervisor host by using the Sconfig.cmd tool.

Issue	Action to Take
You may notice poor network performance if you turn on virtual machine queue (VMQ) for a Hyper-V host that's using a Broadcom network adapter.	For information about a workaround, see the Microsoft documentation, see Poor network performance on virtual machines on a Windows Server 2012 Hyper-V host if VMQ is turned on.

Troubleshooting: Amazon EC2 gateway issues

In the following sections, you can find typical issues that you might encounter working with your gateway deployed on Amazon EC2. For more information about the difference between an on-premises gateway and a gateway deployed in Amazon EC2, see Deploy a default Amazon EC2 host for FSx File Gateway.

Topics

Your gateway activation hasn't occurred after a few moments

Check the following in the Amazon EC2 console:

- Port 80 is open in the security group that you associated with the instance. For more information about adding a security group rule, see Adding a security group rule in the *Amazon EC2 User Guide*.
- The gateway instance is marked as running. In the Amazon EC2 console, the **State** value for the instance should be RUNNING.
- Make sure that your Amazon EC2 instance type meets the minimum requirements, as described in Storage requirements.

After correcting the problem, try activating the gateway again. To do this, open the Storage Gateway console, choose **Deploy a new Gateway on Amazon EC2**, and re-enter the IP address of the instance.

You can't find your EC2 gateway instance in the instance list

If you didn't give your instance a resource tag and you have many instances running, it can be hard to tell which instance you launched. In this case, you can take the following actions to find the gateway instance:

- Check the name of the Amazon Machine Image (AMI) on the **Description** tab of the instance. An instance based on the Storage Gateway AMI should start with the text **`aws-storage-gateway-ami`**.
- If you have several instances based on the Storage Gateway AMI, check the instance launch time to find the correct instance.

You want to connect to your gateway instance using the Amazon EC2 serial console

You can use the Amazon EC2 serial console to troubleshoot boot, network configuration, and other issues. For instructions and troubleshooting tips, see Amazon EC2 Serial Console in the *Amazon Elastic Compute Cloud User Guide*.

You want Support to help troubleshoot your Amazon EC2 gateway

Storage Gateway provides a local console you can use to perform several maintenance tasks, including allowing Support to access your gateway to assist you with troubleshooting gateway issues. By default, Support access to your gateway is turned off. You turn on this access through the Amazon EC2 local console. You log in to the Amazon EC2 local console through a Secure Shell (SSH). To successfully log in through SSH, your instance's security group must have a rule that opens TCP port 22.

> **Note**
>
> If you add a new rule to an existing security group, the new rule applies to all instances that use that security group. For more information about security groups and how to add a security group rule, see Amazon EC2 security groups in the *Amazon EC2 User Guide*.

To let Support connect to your gateway, you first log in to the local console for the Amazon EC2 instance, navigate to the Storage Gateway's console, and then provide the access.

To turn on Support access for a gateway deployed on an Amazon EC2 instance

1. Log in to the local console for your Amazon EC2 instance. For instructions, go to Connect to your instance in the *Amazon EC2 User Guide*.

   You can use the following command to log in to the EC2 instance's local console.

   ```
   ssh -i PRIVATE-KEY admin@INSTANCE-PUBLIC-DNS-NAME
   ```

   > **Note**
   >
   > The *PRIVATE-KEY* is the `.pem` file containing the private certificate of the EC2 key pair that you used to launch the Amazon EC2 instance. For more information, see Retrieving the public key for your key pair in the *Amazon EC2 User Guide*.
   > The *INSTANCE-PUBLIC-DNS-NAME* is the public Domain Name System (DNS) name of your Amazon EC2 instance that your gateway is running on. You obtain this public DNS name by selecting the Amazon EC2 instance in the EC2 console and clicking the **Description** tab.

2. At the prompt, enter **`6 - Command Prompt`** to open the Support Channel console.
3. Enter **h** to open the **AVAILABLE COMMANDS** window.
4. Do one of the following:

   - If your gateway is using a public endpoint, in the **AVAILABLE COMMANDS** window, enter **`open-support-channel`** to connect to customer support for Storage Gateway. Allow TCP port 22 so you can open a support channel to AWS. When you connect to customer support, Storage Gateway assigns you a support number. Make a note of your support number.
   - If your gateway is using a VPC endpoint, in the **AVAILABLE COMMANDS** window, enter **`open-support-channel`**. If your gateway is not activated, provide the VPC endpoint or IP address to connect to customer support for Storage Gateway. Allow TCP port 22 so you can open a support channel to AWS. When you connect to customer support, Storage Gateway assigns you a support number. Make a note of your support number.

> **Note**
>
> The channel number is not a Transmission Control Protocol/User Datagram Protocol (TCP/UDP) port number. Instead, the gateway makes a Secure Shell (SSH) (TCP 22) connection to Storage Gateway servers and provides the support channel for the connection.

5. After the support channel is established, provide your support service number to Support so Support can provide troubleshooting assistance.
6. When the support session is completed, enter **q** to end it. Don't close the session until Amazon Web Services Support notifies you that the support session is complete.
7. Enter **`exit`** to exit the Storage Gateway console.
8. Follow the console menus to log out of the Storage Gateway instance.

Troubleshooting: hardware appliance issues

> **Note**
>
> End of availability notice: As of May 12, 2025, the AWS Storage Gateway Hardware Appliance will no longer be offered. Existing customers with the AWS Storage Gateway Hardware Appliance can continue to use and receive support until May 2028. As an alternative, you can use the AWS Storage Gateway service to give your applications on-premises and in-cloud access to virtually unlimited cloud storage.

The following topics discuss issues that you might encounter with the AWS Storage Gateway Hardware Appliance, and suggestions on troubleshooting these.

Topics

- You can't find the hardware appliance serial number
- Where to obtain hardware appliance support

You can't determine the service IP address

When attempting to connect to your service, make sure that you are using the service's IP address and not the host IP address. Configure the service IP address in the service console, and the host IP address in the hardware console. You see the hardware console when you start the hardware appliance. To go to the service console from the hardware console, choose **Open Service Console**.

How do you perform a factory reset?

If you need to perform a factory reset on your appliance, contact the AWS Storage Gateway Hardware Appliance team for support, as described in the Support section following.

How do you perform a remote restart?

If you need to perform a remote restart of your appliance, you can do so using the Dell iDRAC management interface. For more information, see iDRAC9 Virtual Power Cycle: Remotely power cycle Dell EMC PowerEdge Servers on the Dell Technologies InfoHub website.

Where do you obtain Dell iDRAC support?

The Dell PowerEdge server comes with the Dell iDRAC management interface. We recommend the following:

- If you use the iDRAC management interface, you should change the default password. For more information about the iDRAC credentials, see Dell PowerEdge - What is the default sign-in credentials for iDRAC?.
- Make sure that the firmware is up-to-date to prevent security breaches.
- Moving the iDRAC network interface to a normal (`em`) port can cause performance issues or prevent the normal functioning of the appliance.

You can't find the hardware appliance serial number

You can find the serial number for your AWS Storage Gateway Hardware Appliance using the Storage Gateway console.

To find the hardware appliance serial number:

1. Open the Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. Choose **Hardware** from the navigation menu on the left side of the page.
3. Select your hardware appliance from the list.
4. Locate the **Serial Number** field on the **Details** tab for your appliance.

Where to obtain hardware appliance support

To contact AWS about technical support for your hardware appliance, see Support.

The Support team might ask you to activate the support channel to troubleshoot your gateway issues remotely. You don't need this port to be open for the normal operation of your gateway, but it is required for troubleshooting. You can activate the support channel from the hardware console as shown in the procedure following.

To open a support channel for AWS

1. Open the hardware console.
2. Choose **Open Support Channel** at the bottom of the main page of the hardware console, and then press `Enter`.

   The assigned port number should appear within 30 seconds if there are no network connectivity or firewall issues. For example:

   Status: Open on port 19599
3. Note the port number and provide it to Support.

Troubleshooting: File Gateway issues

You can configure your File Gateway to write log entries to a Amazon CloudWatch log group. If you do, you receive notifications about gateway health status and about any errors that the gateway encounters. You can find information about these error and health notifications in CloudWatch Logs.

In the following sections, you can find information that can help you understand the cause of each error and health notification and how to fix issues.

Topics

Error: FileMissing

The `FileMissing` error is similar to the `ObjectMissing` error, and the steps to resolve it are identical. You can get a `FileMissing` error when a writer other than the specified File Gateway deletes the specified file from the Amazon FSx. Any subsequent uploads to Amazon FSx or retrievals from Amazon FSx for the object fail.

To resolve a FileMissing error

1. Save the latest copy of the file to the local file system of your SMB client (you need this file copy in step 3).
2. Delete the file from the File Gateway using your SMB client.
3. Copy the latest version of the file that you saved in step 1 Amazon FSx using your SMB client. Do this through your File Gateway.

Error: FsxFileSystemAuthenticationFailure

You can get an `FsxFileSystemAuthenticationFailure` error when the credentials provided while attaching the filesystem expired or, its privileges have been revoked.

To resolve an FsxFileSystemAuthenticationFailure error

1. Ensure that the credentials provided at the time of attaching the Amazon FSx file system are still valid.
2. Ensure that the user has all necessary permissions as described in Attach an Amazon FSx for Windows File Server file system.

Error: FsxFileSystemConnectionFailure

You can get an `FsxFileSystemConnectionFailure` error when the Amazon FSx server is inaccessible from the gateway machine.

To resolve an FsxFileSystemConnectionFailure error

1. Ensure that all the firewall and VPC rules are allowing the connection between the gateway machine and the Amazon FSx server.
2. Ensure that the Amazon FSx server is running.

Error: FsxFileSystemFull

You can get an `FsxFileSystemFull` error when there is not enough free disk space in the Amazon FSx file system.

To resolve an FsxFileSystemFull error

- Increase the storage space for the Amazon FSx file system.

Error: GatewayClockOutOfSync

You can get a `GatewayClockOutOfSync` error when the gateway detects a difference of 5 minutes or more between the local system time and the time reported by the AWS Storage Gateway servers. Clock synchronization issues can negatively impact connectivity between the

gateway and AWS. If the gateway clock is out of sync, I/O errors might occur for NFS and SMB connections, and SMB users might experience authentication errors.

To resolve a GatewayClockOutOfSync error

- Check the network configuration between the gateway and the NTP server. For more information about synchronizing the gateway VM time and updating the NTP server configuration, see Configuring a Network Time Protocol (NTP) server for your gateway.

Error: InvalidFileState

You can get an `InvalidFileState` error when a writer other than the specified gateway modifies the specified file in the specified file share. As a result, the state of the file on the gateway doesn't match its state in Amazon FSx. Any subsequent uploads or retrievals of the file from Amazon FSx could fail.

To resolve an InvalidFileState error

1. Save the latest copy of the file to the local file system of your SMB client (you need this file to copy in step 4). If the version of the file in Amazon FSx is the latest, download that version. You can do this by directly accessing the Amazon FSx share using any SMB client.
2. Delete the file in Amazon FSx directly.
3. Delete the file from the gateway using your SMB client.
4. Using your SMB client, copy the latest version of the file that you saved in step 1, *through your File Gateway*,to Amazon FSx.

Error: ObjectMissing

You can get an `ObjectMissing` error when a writer other than the specified File Gateway deletes the specified file from the Amazon FSx. Any subsequent uploads to Amazon FSx or retrievals from Amazon FSx for the object fail.

To resolve an ObjectMissing error

1. Save the latest copy of the file to the local file system of your SMB client (you need this file copy in step 3).
2. Delete the file from the File Gateway using your SMB client.

3. Copy the latest version of the file that you saved in step 1 Amazon FSx using your SMB client. Do this through your File Gateway.

Error: DroppedNotifications

You might see a `DroppedNotifications` error instead of other expected types of CloudWatch log entries when free storage space on your gateway's root disk is less than 1 GB, or if more than 100 health notifications are generated within a 1 minute interval. In these circumstances, the gateway stops generating detailed CloudWatch log notifications as a precautionary measure.

To resolve a DroppedNotifications error

1. Check the `Root Disk Usage` metric on the **Monitoring** tab for your gateway in the Storage Gateway console to determine whether available root disk space is running low.
2. Increase the size of the gateway's root storage disk if available space is less than 1 GB. Refer to your virtual machine hypervisor's documentation for instructions.

   To increase root disk size for Amazon EC2 gateways, see Request modifications to your EBS volumes in the *Amazon Elastic Compute Cloud User Guide*.

   > **Note**
   >
   > It is not possible to increase the root disk size for the AWS Storage Gateway Hardware Appliance.

3. Restart your gateway.

Notification: HardReboot

You can get a `HardReboot` notification when the gateway VM is restarted unexpectedly. Such a restart can be due to loss of power, a hardware failure, or another event. For VMware gateways, a reset by vSphere High Availability Application Monitoring can cause this event.

When your gateway runs in such an environment, check for the presence of the `HealthCheckFailure` notification and consult the VMware events log for the VM.

Notification: Reboot

You can get a reboot notification when the gateway VM is restarted. You can restart a gateway VM by using the VM Hypervisor Management console or the Storage Gateway console. You can also restart by using the gateway software during the gateway's maintenance cycle.

If the time of the reboot is within 10 minutes of the gateway's configured maintenance start time, this reboot is probably a normal occurrence and not a sign of any problem. If the reboot occurred significantly outside the maintenance window, check whether the gateway was restarted manually.

Troubleshooting: Active Directory domain issues

FSx File Gateway doesn't generate specific log messages for Active Directory domain issues. If you have trouble joining your gateway to your Active Directory domain, do the following:

- Verify that the gateway is not attempting to use a read-only domain controller (RODC) to join the domain.
- Verify that the gateway is configured to use the correct DNS servers.

  For example, if you are trying to join an Amazon EC2 gateway instance to an AWS-managed Active Directory, verify that the DHCP option set for your EC2 VPC specifies the AWS-managed Active Directory DNS servers.

  DNS servers that you configure through the VPC DHCP options set are provided to the all EC2 instances in the VPC. If you want to specify a DNS server for an individual gateway, you can do so using that gateway's EC2 local console.

  For on-premises gateways, you specify a DNS server using the VM local console.
- Verify gateway network connectivity by running the following commands from the command prompt in the gateway's local console. Replace the highlighted variables with the actual domain name and IP addresses from your deployment.

  ```
  dig -d ExampleDomainName
  ncport -d ExampleDomainControllerIPAddress -p 445
  ncport -d ExampleDomainControllerIPAddress -p 389
  ```
- Verify that your Active Directory service account has the requisite permissions. For more information, see Active Directory service account permission requirements.
- Verify that the gateway joins the correct Organizational Unit (OU).

Joining a domain creates an Active Directory computer account in the default computers container (which is not an OU), using the gateway's **Gateway ID** as the account name (for example, SGW-1234ADE). It is not possible to customize the name of this account.

If your Active Directory environment has a designated OU for new computer objects, you must specify that OU when joining the domain.

If you encounter access denied errors when attempting to join the designated OU, check with your Active Directory domain administrator. The administrator may need to pre-stage the gateway's computer account before it can join the domain. For more information, see How can I troubleshoot issues with joining my Storage Gateway file gateway to a domain for Microsoft Active Directory authentication?.

- Verify that your gateway's hostname is resolvable in DNS by running the following command from the command prompt in the gateway's local console. Replace the highlighted variable with the actual hostname for your gateway.

```
dig -d ExampleHostName -r A
```

If you configured a custom hostname for your gateway, you must manually add a DNS A-record that points to its IP address.

- Verify that network latency between the gateway and the domain controller is reasonably low. The query to join a domain can time out if the gateway does not receive a response from the domain controller within 20 seconds.

If you join the gateway to the domain using the JoinDomain CLI command, you can can add the `--timeout-in-seconds` flag to increase the timeout to a maximum of 3,600 seconds.

- Verify that the Active Directory user you are using to join the gateway to the domain has the privileges required to do so.

Troubleshooting: Using CloudWatch metrics

You can find information following about actions to address issues using Amazon CloudWatch metrics with Storage Gateway.

Topics

- Your gateway reacts slowly when browsing directories

- Your gateway isn't responding
- You do not see files in your Amazon FSx file system
- You do not see older snapshots in your Amazon FSx file system
- Your gateway is slow transferring data to Amazon FSx
- Your gateway backup job fails or there are errors when writing to your gateway

Your gateway reacts slowly when browsing directories

If your File Gateway reacts slowly when you run the **ls** command or browse directories, check the `IndexFetch` and `IndexEviction` CloudWatch metrics:

- If the `IndexFetch` metric is greater than 0 when you run an `ls` command or browse directories, your File Gateway started without information on the contents of the directory affected and had to access FSx for Windows File Server. Subsequent efforts to list the contents of that directory should go faster.
- If the `IndexEviction` metric is greater than 0, it means that your File Gateway has reached the limit of what it can manage in its cache at that time. In this case, your File Gateway has to free some storage space from the least recently accessed directory to list a new directory. If this occurs frequently and there is a performance impact, contact Support.

  Discuss with Support the contents of the related Amazon FSx file system and recommendations to improve performance based on your use case.

Your gateway isn't responding

If your File Gateway isn't responding, do the following:

- If there was a recent reboot or software update, then check the `IOWaitPercent` metric. This metric shows the percentage of time that the CPU is idle when there is an outstanding disk I/O request. In some cases, this might be high (10 or greater) and might have risen after the server was rebooted or updated. In these cases, then your File Gateway might be bottlenecked by a slow root disk as it rebuilds the index cache to RAM. You can address this issue by using a faster physical disk for the root disk.
- If the `MemUsedBytes` metric is at or nearly the same as the `MemTotalBytes` metric, then your File Gateway is running out of available RAM. Make sure that your File Gateway has at least

the minimum required RAM. If it already does, consider adding more RAM to your File Gateway based on your workload and use case.

If the file share is SMB, the issue might also be due to the number of SMB clients connected to the file share. To see the number of clients connected at any given time, check the `SMBV(1/2/3)Sessions` metric. If there are many clients connected, you might need to add more RAM to your File Gateway.

You do not see files in your Amazon FSx file system

If you notice that files on the gateway are not reflected in the Amazon FSx file system, check the `FilesFailingUpload` metric. If the metric reports that some files are failing upload, check your health notifications. When files fail to upload, the gateway generates a health notification containing more details on the issue.

You do not see older snapshots in your Amazon FSx file system

Some file operations on the FSx File Gateway, such as top-level folder renames or permission changes, can result in multiple file operations that lead to a high I/O load on your FSx for Windows File Server file system. If your file system doesn't have enough performance resources for your workload, the file system might delete shadow copies because it prioritizes availability for ongoing I/O over historical shadow copy retention.

In the Amazon FSx console, check the **Monitoring and performance** page to see if your file system is under-provisioned. If it is, you can switch to SSD storage, increase throughput capacity, or increase SSD IOPS to handle your workload.

Your gateway is slow transferring data to Amazon FSx

If your File Gateway is slow transferring data to Amazon FSx for Windows File Server, do the following:

- If the `CachePercentDirty` metric is 80 or greater, your File Gateway is writing data faster to disk than it can upload the data to Amazon FSx for Windows File Server. Consider increasing the bandwidth for upload from your File Gateway, adding one or more cache disks, or slowing down client writes, or increase the throughput capacity for associated Amazon FSx for Windows File Server.
- If the `CachePercentDirty` metric is low, check the `IoWaitPercent` metric. If `IoWaitPercent` is greater than 10, your File Gateway might be bottlenecked by the speed of

the local cache disk. We recommend local solid state drive (SSD) disks for your cache, preferably NVM Express (NVMe). If such disks aren't available, try using multiple cache disks from separate physical disks for a performance improvement.

Your gateway backup job fails or there are errors when writing to your gateway

If your File Gateway backup job fails or there are errors when writing to your File Gateway, do the following:

- If the `CachePercentDirty` metric is 90 percent or greater, your File Gateway can't accept new writes to disk because there is not enough available space on the cache disk. To see how fast your File Gateway is uploading to FSx for Windows File Server, view the `CloudBytesUploaded` metric. Compare that metric with the `WriteBytes` metric, which shows how fast the client is writing files to your File Gateway. If the SMB client is writing to your File Gateway faster than it can upload to FSx for Windows File Server, add more cache disks to cover the size of the backup job at a minimum. Or, increase the upload bandwidth.
- If a large file copy such as backup job fails but the `CachePercentDirty` metric is less than 80 percent, your File Gateway might be hitting a client-side session timeout. For SMB, you can increase this timeout using the PowerShell command `Set-SmbClientConfiguration -SessionTimeout 300`. Running this command sets the timeout to 300 seconds.

High Availability Health Notifications

When running your gateway on the VMware vSphere High Availability (HA) platform, you may receive health notifications. For more information about health notifications, see Troubleshooting: high availability issues.

Troubleshooting: high availability issues

You can find information following about actions to take if you experience availability issues.

Topics

Health notifications

When you run your gateway on VMware vSphere HA, all gateways produce the following health notifications to your configured Amazon CloudWatch log group. These notifications go into a log stream called `AvailabilityMonitor`.

Topics

Notification: Reboot

You can get a reboot notification when the gateway VM is restarted. You can restart a gateway VM by using the VM Hypervisor Management console or the Storage Gateway console. You can also restart by using the gateway software during the gateway's maintenance cycle.

Action to Take

If the time of the reboot is within 10 minutes of the gateway's configured maintenance start time, this is probably a normal occurrence and not a sign of any problem. If the reboot occurred significantly outside the maintenance window, check whether the gateway was restarted manually.

Notification: HardReboot

You can get a `HardReboot` notification when the gateway VM is restarted unexpectedly. Such a restart can be due to loss of power, a hardware failure, or another event. For VMware gateways, a reset by vSphere High Availability Application Monitoring can cause this event.

Action to Take

When your gateway runs in such an environment, check for the presence of the `HealthCheckFailure` notification and consult the VMware events log for the VM.

Notification: HealthCheckFailure

For a gateway on VMware vSphere HA, you can get a `HealthCheckFailure` notification when a health check fails and a VM restart is requested. This event also occurs during a test to

monitor availability, indicated by an `AvailabilityMonitorTest` notification. In this case, the `HealthCheckFailure` notification is expected.

> **Note**
>
> This notification is for VMware gateways only.

Action to Take

If this event repeatedly occurs without an `AvailabilityMonitorTest` notification, check your VM infrastructure for issues (storage, memory, and so on). If you need additional assistance, contact Support.

Notification: AvailabilityMonitorTest

For a gateway on VMware vSphere HA, you can get an `AvailabilityMonitorTest` notification when you run a test of the Availability and application monitoring system in VMware.

Metrics

The `AvailabilityNotifications` metric is available on all gateways. This metric is a count of the number of availability-related health notifications generated by the gateway. Use the `Sum` statistic to observe whether the gateway is experiencing any availability-related events. Consult with your configured CloudWatch log group for details about the events.

Best practices for File Gateway

This section contains the following topics, which provide information about the best practices for working with gateways, file shares, buckets, and data. We recommend that you familiarize yourself with the information outlined in this section, and attempt to follow these guidelines in order to avoid problems with your AWS Storage Gateway. For additional guidance on diagnosing and solving common issues you might encounter with your deployment, see Troubleshooting problems with your Storage Gateway deployment.

Topics

Best practices: recovering your data

Although it is rare, your gateway might encounter an unrecoverable failure. Such a failure can occur in your virtual machine (VM), the gateway itself, the local storage, or elsewhere. If a failure occurs, we recommend that you follow the instructions in the appropriate section following to recover your data.

> **Important**
>
> Storage Gateway doesn't support recovering a gateway VM from a snapshot that is created by your hypervisor or from your Amazon EC2 Amazon Machine Image (AMI). If your gateway VM malfunctions, activate a new gateway and recover your data to that gateway using the instructions following.

Recovering from an unexpected virtual machine shutdown

If your VM shuts down unexpectedly, for example during a power outage, your gateway becomes unreachable. When power and network connectivity are restored, your gateway becomes reachable and starts to function normally. Following are some steps you can take at that point to help recover your data:

- If an outage causes network connectivity issues, you can troubleshoot the issue. For information about how to test network connectivity, see Testing your gateway's network connectivity.

Recovering your data from a malfunctioning cache disk

If your cache disk encounters a failure, we recommend you use the following steps to recover your data depending on your situation:

- If the malfunction occurred because a cache disk was removed from your host, shut down the gateway, re-add the disk, and restart the gateway.

Recovering your data from an inaccessible data center

If your gateway or data center becomes inaccessible for some reason, you can recover your data to another gateway in a different data center or recover to a gateway hosted on an Amazon EC2 instance. If you don't have access to another data center, we recommend creating the gateway on an Amazon EC2 instance. The steps you follow depends on the gateway type you are covering the data from.

To recover data from a File Gateway in an inaccessible data center

For File Gateway, you map a new file system to the FSx for Windows File Server that contains the data you want to recover.

1. Create and activate a new File Gateway on an Amazon EC2 host. For more information, see Deploy a default Amazon EC2 host for FSx File Gateway.
2. Create a new file system on the EC2 gateway you created. For more information, see Create an FSx for Windows File Server file system.
3. Mount your file system on your client and map it to the FSx for Windows File Server that contains the data that you want to recover. For more information, see Mount and use your file share.

Restoring from backups or snapshots directly on Amazon FSx

In some cases, you might need to restore data on your Amazon FSx file system directly, using a backup or snapshot from an earlier point in time. In these instances, there is a risk of creating a dual-writer scenario between the backup application and the FSx File Gateway, which can result in

stuck or mis-matched files. To avoid problems when restoring your Amazon FSx file system from backups or snapshots, use the following procedure.

> **Note**
>
> Any cached data currently stored on your FSx File Gateway will not be valid after you restore your Amazon FSx file system from a backup or snapshot using this procedure.

To avoid problems when restoring your Amazon FSx file system from backups or snapshots

1. Detach the Amazon FSx file system from the FSx File Gateway using the Storage Gateway console.
2. Restore the backup or snapshot directly on your Amazon FSx file system.
3. Reattach the Amazon FSx file system to the FSx File Gateway using the Storage Gateway console.

Clean up unnecessary resources

As a best practice, we recommend cleaning up Storage Gateway resources to avoid unexpected or unnecessary charges. For example, if you created a gateway as a demonstration exercise or a test, consider deleting it and its virtual appliance from your deployment. Use the following procedure to clean up resources.

To clean up resources you don't need

1. If you no longer plan to continue using a gateway, delete it. For more information, see Deleting your gateway and removing associated resources.
2. Delete the Storage Gateway VM from your on-premises host. If you created your gateway on an Amazon EC2 instance, terminate the instance.

Additional Storage Gateway resources

This section contains the following topics, which provide additional information and resources related to setting up and using AWS Storage Gateway:

Topics

- Host setup - Learn how to deploy and configure a virtual machine host for your gateway.
- Using Storage Gateway with VMware HA - Learn how to set up Storage Gateway to work with VMware vSphere high availability features.
- Getting activation key - Learn where to find the activation key that you need to provide when you deploy a new gateway.
- Using AWS Direct Connect - Learn how to create a dedicated network connection between your on-premises gateway and the AWS cloud.
- Active Directory permissions - Learn which permissions your service account must have to be able to join your gateway to your Active Directory domain.
- Getting the IP address for your gateway appliance - Learn where to find the gateway's virtual machine host IP address, which you need to provide when you deploy a new gateway.
- Understanding resources and resource IDs - Learn how AWS identifies the resources and subresources that are created by Storage Gateway.
- Tagging your resources - Learn how to use metadata tags to categorize your resources and make them easier to manage.
- Open-source components - Learn about the third-party tools and licenses that are used to deliver Storage Gateway functionality.
- Quotas - Learn about limits and quotas for File Gateway, including minimum and maximum limitations for file shares and local cache disks.

Deploying and configuring the gateway VM host

The following topics provide information about setting up the virtual machine host platform for your gateway.

Topics

- Deploy a default Amazon EC2 host for FSx File Gateway

- Deploy a customized Amazon EC2 host for FSx File Gateway
- Modify Amazon EC2 instance metadata options
- Synchronize VM time with Hyper-V or Linux KVM host time
- Synchronize VM time with VMware host time
- Configuring network adapters for your gateway
- Using VMware vSphere High Availability with Storage Gateway

Deploy a default Amazon EC2 host for FSx File Gateway

This topic lists the steps to deploy an Amazon EC2 host using the default specifications.

You can deploy and activate an Amazon FSx File Gateway on an Amazon Elastic Compute Cloud (Amazon EC2) instance. The AWS Storage Gateway Amazon Machine Image (AMI) is available as a community AMI.

> **Note**
>
> Storage Gateway community AMIs are published and fully supported by AWS. You can see that the publisher is AWS, a verified provider.

1. To set up the Amazon EC2 instance, choose **Amazon EC2** as the **Host platform** in the **Platform options** section of the workflow. For instructions on configuring the Amazon EC2 instance, see Deploying an Amazon EC2 instance to host your Amazon FSx File Gateway.
2. Select **Launch instance** to open the AWS Storage Gateway AMI template in the Amazon EC2 console and customize additional settings such as **Instance types**, **Network settings** and **Configure storage**.
3. Optionally, you can select **Use default settings** in the Storage Gateway console to deploy an Amazon EC2 instance with the default configuration.

   The Amazon EC2 instance that **Use default settings** creates has the following default specifications:

   - **Instance type** — *m5.xlarge*
   - **Network Settings**
     - For **VPC**, select the VPC that you want your EC2 instance to run in.

- For **Subnet**, specify the subnet that your EC2 instance should be launched in.

> **Note**
>
> VPC subnets will appear in the drop down only if they have the auto-assign public IPv4 address setting activated from the VPC management console.

- **Auto-assign Public IP** — *Activated*
- An EC2 security group is created and associated with the EC2 Instance. The security group has the following inbound port rules:

> **Note**
>
> You will need Port 80 open during gateway activation. The port is closed immediately following activation. Thereafter, your EC2 instance can only be accessed over the other ports from the selected VPC.
> The file shares on your gateway are only accessible from the hosts in the same VPC as the gateway. If the file shares need to be accessed from hosts outside of the VPC, you should update the appropriate security group rules.
> You can edit security groups at any time by navigating to the Amazon EC2 instance details page, selecting **Security**, navigating to **Security group details**, and choosing the security group ID.

Port	Protocol	File System Protocol
80	TCP	HTTP access for activation
137	UDP	NetBIOS
138	UDP	NetBIOS
139	TCP, UDP	SMB

Port	Protocol	File System Protocol
389	TCP	LDAP
445	TCP	SMB

- **Configure storage**

Default Settings	AMI Root Volume	Volume 2 Cache
Device Name		'/dev/sdb'
Size	80 Gib	165 GiB
Volume Type	gp3	gp3
IOPS	3000	3000
Delete on termination	Yes	Yes
Encrypted	No	No
Throughput	125	125

Deploy a customized Amazon EC2 host for FSx File Gateway

You can deploy and activate an Amazon FSx File Gateway on an Amazon Elastic Compute Cloud (Amazon EC2) instance. The AWS Storage Gateway Amazon Machine Image (AMI) is available as a community AMI.

> **Note**
>
> Storage Gateway community AMIs are published and fully supported by AWS. You can see that the publisher is AWS, a verified provider.
> FSx File Gateway AMIs use the following naming convention. The version number appended to the AMI name changes with each version release.
> `aws-storage-gateway-FILE_FSX_SMB-2.2.3`

To deploy an Amazon EC2 instance to host your Amazon FSx File Gateway

1. Start setting up a new gateway using the Storage Gateway console. For instructions, see Set up an Amazon FSx File Gateway. When you reach the **Platform options** section, choose **Amazon EC2** as the **Host platform**, then use the following steps to launch the Amazon EC2 instance that will host your File Gateway.
2. Choose **Launch instance** to open the AWS Storage Gateway AMI template in the Amazon EC2 console, where you can configure additional settings.

   Use **Quicklaunch** to launch the Amazon EC2 instance with default settings. For more information on Amazon EC2 Quicklaunch default specifications, see Quicklaunch Configuration Specifications for Amazon EC2.
3. For **Name**, enter a name for the Amazon EC2 instance. After the instance is deployed, you can search for this name to find your instance on list pages in the Amazon EC2 console.
4. In the **Instance type** section, for **Instance type**, choose the hardware configuration for your instance. The hardware configuration must meet certain minimum requirements to support your gateway. We recommend starting with the **m5.xlarge** instance type, which meets the minimum hardware requirements for your gateway to function properly. For more information, see Requirements for Amazon EC2 instance types.

   You can resize your instance after you launch, if necessary. For more information, see Resizing your instance in the *Amazon EC2 User Guide*.

   > **Note**
   >
   > Certain instance types, particularly i3 EC2, use NVMe SSD disks. These can cause problems when you start or stop File Gateway; for example, you can lose data from the cache. Monitor the `CachePercentDirty` Amazon CloudWatch metric, and only start

or stop your system when that parameter is `0`. To learn more about monitoring metrics for your gateway, see Storage Gateway metrics and dimensions in the CloudWatch documentation.

5. In the **Key pair (login)** section, for **Key pair name - *required***, select the key pair you want to use to securely connect to your instance. You can create a new key pair if necessary. For more information, see Create a key pair in the *Amazon Elastic Compute Cloud User Guide for Linux Instances*.
6. In the **Network settings** section, review the preconfigured settings and choose **Edit** to make changes to the following fields:
   a. For **VPC - *required***, choose the VPC where you want to launch your Amazon EC2 instance. For more information, see How Amazon VPC works in the *Amazon Virtual Private Cloud User Guide*.
   b. (Optional) For **Subnet**, choose the subnet where you want to launch your Amazon EC2 instance.
   c. For **Auto-assign Public IP**, choose **Enable**.
7. In the **Firewall (security groups)** subsection, review the preconfigured settings. You can change the default name and description of the new security group to be created for your Amazon EC2 instance if you want, or choose to apply firewall rules from an existing security group instead.
8. In the **Inbound security groups rules** subsection, add firewall rules to open the ports that clients will use to connect to your instance. For more information on the ports required for Amazon FSx File Gateway, see Port requirements. For more information on adding firewall rules, see Security group rules in the *Amazon Elastic Compute Cloud User Guide for Linux Instances*.

   Note

   Amazon FSx File Gateway requires TCP port 80 to be open for inbound traffic and one-time HTTP access during gateway activation. After activation, you can close this port. Additionally, you must open TCP port 445 for SMB access, UDP port 137 for NetBIOS access, UDP port 138 for NetBIOS access, and TCP port 389 for LDAP access.

9. In the **Advanced network configuration** subsection, review the preconfigured settings and make changes if necessary.

10. In the **Configure storage** section, choose **Add new volume** to add storage to your gateway instance.

    > **Important**
    >
    > You must add at least one Amazon EBS volume with at least **150 GiB** capacity for cache storage in addition to the preconfigured **Root volume**. For increased performance, we recommend allocating multiple EBS volumes for cache storage with at least 150 GiB each.

11. In the **Advanced details** section, review the preconfigured settings and make changes if necessary.
12. Choose **Launch instance** to launch your new Amazon EC2 gateway instance with the configured settings.
13. To verify that your new instance launched successfully, navigate to the **Instances** page in the Amazon EC2 console and search for your new instance by name. Ensure that that **Instance state** displays **Running** *with a green check mark*, and that the **Status check** is complete, and *shows a green check mark*.
14. Select your instance from the details page. Copy the **Public IPv4 address** from the **Instance summary** section, then return to the **Set up gateway** page in the Storage Gateway console to resume setting up your Amazon FSx File Gateway.

You can determine the AMI ID to use for launching a File Gateway by using the Storage Gateway console or by querying the AWS Systems Manager parameter store.

To determine the AMI ID, do one of the following:

- Start setting up a new gateway using the Storage Gateway console. For instructions, see Set up an Amazon FSx File Gateway. When you reach the **Platform options** section, choose **Amazon EC2** as the **Host platform**, then choose **Launch instance** to open the AWS Storage Gateway AMI template in the Amazon EC2 console.

  You are redirected to the EC2 community AMI page, where you can see the AMI ID for your AWS Region in the URL.

- Query the Systems Manager parameter store. You can use the AWS CLI or Storage Gateway API to query the Systems Manager public parameter under the namespace `/aws/service/`

`storagegateway/ami/FILE_FSX_SMB/latest`. For example, using the following CLI command returns the ID of the current AMI in the AWS Region you specify.

```
aws --region us-east-2 ssm get-parameter --name /aws/service/storagegateway/ami/
FILE_FSX_SMB/latest
```

The CLI command returns output similar to the following.

```
{
    "Parameter": {
        "Type": "String",
        "LastModifiedDate": 1561054105.083,
        "Version": 18,
         "ARN": "arn:aws:ssm:us-east-2::parameter/aws/service/storagegateway/ami/
FILE_FSX_SMB/latest",
         "Name": "/aws/service/storagegateway/ami/FILE_FSX_SMB/latest",,
        "Value": "ami-033d1edba5606cffb"
    }
}
```

Modify Amazon EC2 instance metadata options

The instance metadata service (IMDS) is an on-instance component that provides secure access to Amazon EC2 instance metadata. An instance can be configured to accept incoming metadata requests that use IMDS Version 1 (IMDSv1) or require that all metadata requests use IMDS Version 2 (IMDSv2). IMDSv2 uses session-oriented requests and mitigates several types of vulnerabilities that could be used to try to access the IMDS. For information about IMDSv2, see How Instance Metadata Service Version 2 works in the *Amazon Elastic Compute Cloud User Guide*.

We recommend that you require IMDSv2 for all Amazon EC2 instances that host Storage Gateway. IMDSv2 is required by default on all newly launched gateway instances. If you have existing instances that are still configured to accept IMDSv1 metadata requests, see Require the use of IMDSv2 in the *Amazon Elastic Compute Cloud User Guide* for instructions to modify your instance metadata options to require the use of IMDSv2. Applying this change does not require an instance reboot.

Synchronize VM time with Hyper-V or Linux KVM host time

For a gateway deployed on VMware ESXi, setting the hypervisor host time and synchronizing the virtual machine time to the host is sufficient to avoid time drift. For more information, see Synchronize VM time with VMware host time. For a gateway deployed on Microsoft Hyper-V or Linux KVM, we recommend that you periodically check the virtual machine time using the procedure described following.

To view and synchronize the time of a hypervisor gateway virtual machine to a Network Time Protocol (NTP) server

1. Log in to your gateway's local console:
   - For more information on logging in to the Microsoft Hyper-V local console, see Access the Gateway Local Console with Microsoft Hyper-V.
   - For more information on logging in to the local console for Linux Kernel-based Virtual Machine (KVM), see Accessing the Gateway Local Console with Linux KVM.
2. On the **Storage Gateway Configuration** main menu screen, enter the corresponding numeral to select **System Time Management**.
3. On the **System Time Management** menu screen, enter the corresponding numeral to select **View and Synchronize System Time**.

   The gateway local console displays the current system time and compares it with the time reported by the NTP server, then reports the exact discrepancy between the two times in seconds.
4. If the time discrepancy is greater than 60 seconds, enter **y** to synchronize the system time with NTP time. Otherwise, enter **n**.

   Time synchronization might take a few moments.

Synchronize VM time with VMware host time

To successfully activate your gateway, you must ensure that your VM time is synchronized to the host time, and that the host time is correctly set. In this section, you first synchronize the time on the VM to the host time. Then you check the host time and, if needed, set the host time and configure the host to synchronize its time automatically to a Network Time Protocol (NTP) server.

> **Important**
>
> Synchronizing the VM time with the host time is required for successful gateway activation.

To synchronize VM time with host time

1. Configure your VM time.

   a. In the vSphere client, right-click on the name of your gateway VM in panel on the left side of the application window to open the context menu for the VM, and then choose **Edit Settings**.

      The **Virtual Machine Properties** dialog box opens.

   b. Choose the **Options** tab, and then choose **VMware Tools** from the options list.

   c. Check the **Synchronize guest time with host** option in the **Advanced** section on the right side of the **Virtual Machine Properties** dialog box, and then choose **OK**.

      The VM synchronizes its time with the host.

2. Configure the host time.

   It is important to make sure that your host clock is set to the correct time. If you have not configured your host clock, perform the following steps to set and synchronize it with an NTP server.

   a. In the VMware vSphere client, select the vSphere host node in the left panel, and then choose the **Configuration** tab.

   b. Select **Time Configuration** in the **Software** panel, and then choose the **Properties** link.

      The **Time Configuration** dialog box appears.

   c. Under **Date and Time**, set the date and time for your vSphere host.

   d. Configure the host to synchronize its time automatically to an NTP server.

      i. Choose **Options** in the **Time Configuration** dialog box, and then in the **NTP Daemon (ntpd) Options** dialog box, choose **NTP Settings** in the left panel.

      ii. Choose **Add** to add a new NTP server.

      iii. In the **Add NTP Server** dialog box, type the IP address or the fully qualified domain name of an NTP server, and then choose **OK**.

You can use `pool.ntp.org` as the domain name.

iv. In the **NTP Daemon (ntpd) Options** dialog box, choose **General** in the left panel.

v. Under **Service Commands**, choose **Start** to start the service.

Note that if you change this NTP server reference or add another later, you will need to restart the service to use the new server.

e. Choose **OK** to close the **NTP Daemon (ntpd) Options** dialog box.

f. Choose **OK** to close the **Time Configuration** dialog box.

Configuring network adapters for your gateway

Storage Gateway uses a single VMXNET3 (10 GbE) network adapter by default, but you can configure your gateway to use more than one network adapter so that it can be accessed by multiple IP addresses. You might want to do this in the following situations:

- **Maximizing throughput**– You might want to maximize throughput to a gateway when network adapters are a bottleneck.
- **Application separation** – You might need to separate how your applications write to a gateway's volumes. For example, you might choose to have a critical storage application exclusively use one particular adapter defined for your gateway.
- **Network constraints** – Your application environment might require that you keep your file shares and the initiators that connect to them in an isolated network. This network is different from the network by which the gateway communicates with AWS.

In a typical multiple-adapter use case, one adapter is configured as the route by which the gateway communicates with AWS (that is, as the default gateway). Except for this one adapter, initiators must be in the same subnet as the adapter that contains the file shares to which they connect. Otherwise, communication with the intended targets might not be possible. If a target is configured on the same adapter that is used for communication with AWS, then file share traffic for that target and AWS traffic flows through the same adapter.

In some cases, you might configure one adapter to connect to the Storage Gateway console and then add a second adapter. In such a case, Storage Gateway automatically configures the route table to use the second adapter as the preferred route. For instructions on how to configure multiple adapters, see the following topics:

Topics

Configuring Your Gateway for Multiple NICs on a VMware ESXi Host

The following procedure assumes that your gateway VM already has one network adapter defined, and describes how to add an adapter on VMware ESXi.

To configure your gateway to use an additional network adapter in VMware ESXi host

1. Shut down the gateway.
2. In the VMware vSphere client, select your gateway VM.

   The VM can remain turned on for this procedure.
3. In the client, open the context (right-click) menu for your gateway VM, and choose **Edit Settings**.
4. On the **Hardware** tab of the **Virtual Machine Properties** dialog box, choose **Add** to add a device.
5. Follow the Add Hardware wizard to add a network adapter.

   a. In the **Device Type** pane, choose **Ethernet Adapter** to add an adapter, and then choose **Next**.

   b. In the **Network Type** pane, ensure that **Connect at power on** is selected for **Type**, and then choose **Next**.

      We recommend that you use the VMXNET3 network adapter with Storage Gateway. For more information on the adapter types that might appear in the adapter list, see Network Adapter Types in the ESXi and vCenter Server Documentation.

   c. In the **Ready to Complete** pane, review the information, and then choose **Finish**.
6. Choose the **Summary** tab for the VM, and choose **View All** next to the **IP Address** box. The **Virtual Machine IP Addresses** window displays all the IP addresses you can use to access the gateway. Confirm that a second IP address is listed for the gateway.

> **Note**
>
> It might take several moments for the adapter changes to take effect and the VM summary information to refresh.

7. In the Storage Gateway console, turn on the gateway.
8. In the **Navigation** pane of the Storage Gateway console, choose **Gateways** and choose the gateway to which you added the adapter. Confirm that the second IP address is listed in the **Details** tab.

For information about local console tasks common to VMware, Hyper-V, and KVM hosts, see Performing tasks on the virtual machine local console

Configuring Your Gateway for Multiple NICs in Microsoft Hyper-V Host

The following procedure assumes that your gateway VM already has one network adapter defined and that you are adding a second adapter. This procedure shows how to add an adapter for a Microsoft Hyper-V host.

To configure your gateway to use an additional network adapter in a Microsoft Hyper-V Host

1. On the Storage Gateway console, turn off the gateway.
2. In the Microsoft Hyper-V Manager, select your gateway VM from the **Virtual Machines** panel.
3. If the gateway VM isn't turned off already, right-click the VM name to open the context menu, and then choose **Turn Off**.
4. Right-click the gateway VM name to open the context menu, and then choose **Settings**.
5. In the **Settings** dialog box, under **Hardware**, choose **Add Hardware**.
6. In the **Add Hardware** panel on the right side of the **Settings** dialog box, choose **Network Adapter**, and then choose **Add** to add a device.
7. Configure the network adapter, and then choose **Apply** to apply settings.
8. In the **Settings** dialog box, under **Hardware**, confirm that the new network adapter was added to the hardware list, and then choose **OK**.
9. Turn on the gateway using the Storage Gateway console.

10. In the **Navigation** panel of the Storage Gateway console, choose **Gateways**, then select the gateway to which you added the adapter. Confirm that a second IP address is listed in the **Details** tab.

For information about local console tasks common to VMware, Hyper-V, and KVM hosts, see Performing tasks on the virtual machine local console

Using VMware vSphere High Availability with Storage Gateway

Storage Gateway provides high availability on VMware through a set of application-level health checks integrated with VMware vSphere High Availability (VMware HA). This approach helps protect storage workloads against hardware, hypervisor, or network failures. It also helps protect against software errors, such as connection timeouts and file share or volume unavailability.

With this integration, a gateway deployed in a VMware environment on-premises or in a VMware Cloud on AWS automatically recovers from most service interruptions. It generally does this in under 60 seconds with no data loss.

> **Note**
>
> We recommend doing the following things if you deploy Storage Gateway in a VMware HA cluster:
>
> - Deploy the VMware ESX .ova downloadable package that contains the Storage Gateway VM on only one host in a cluster.
> - When deploying the .ova package, select a data store that is not local to one host. Instead, use a data store that is accessible to all hosts in the cluster. If you select a data store that is local to a host and the host fails, then the data source might not be accessible to other hosts in the cluster and failover to another host might not succeed.
> - With clustering, if you deploy the .ova package to the cluster, select a host when you are prompted to do so. Alternately, you can deploy directly to a host in a cluster.

The following topics describe how to deploy Storage Gateway in a VMware HA cluster:

Topics

- Configure Your vSphere VMware HA Cluster
- Set Up Your Gateway Type

- Deploy the Gateway
- (Optional) Add Override Options for Other VMs on Your Cluster
- Activate Your Gateway
- Test Your VMware High Availability Configuration

Configure Your vSphere VMware HA Cluster

First, if you haven't already created a VMware cluster, create one. For information about how to create a VMware cluster, see Create a vSphere HA Cluster in the VMware documentation.

Next, configure your VMware cluster to work with Storage Gateway.

To configure your VMware cluster

1. On the **Edit Cluster Settings** page in VMware vSphere, make sure that VM monitoring is configured for VM and application monitoring. To do so, set the following values for each option:

   - **Host Failure Response**: **Restart VMs**
   - **Response for Host Isolation**: **Shut down and restart VMs**
   - **Datastore with PDL**: **Disabled**
   - **Datastore with APD**: **Disabled**
   - **VM Monitoring**: **VM and Application Monitoring**
2. Fine-tune the sensitivity of the cluster by adjusting the following values:

   - **Failure interval** – After this interval, the VM is restarted if a VM heartbeat isn't received.
   - **Minimum uptime** – The cluster waits this long after a VM starts to begin monitoring for VM tools' heartbeats.
   - **Maximum per-VM resets** – The cluster restarts the VM a maximum of this many times within the maximum resets time window.
   - **Maximum resets time window** – The window of time in which to count the maximum resets per-VM resets.

   If you aren't sure what values to set, use these example settings:

   - **Failure interval**: **30** seconds

- **Minimum uptime**: **120** seconds
- **Maximum per-VM resets**: **3**
- **Maximum resets time window**: **1** hour

If you have other VMs running on the cluster, you might want to set these values specifically for your VM. You can't do this until you deploy the VM from the .ova. For more information on setting these values, see (Optional) Add Override Options for Other VMs on Your Cluster.

Set Up Your Gateway Type

Use the following procedure to set up the gateway

To download the .ova image for your gateway type

- Download the .ova image for your gateway type from one of the following:
  - File Gateway – Create and activate an Amazon FSx File Gateway

Deploy the Gateway

In your configured cluster, deploy the .ova image to one of the cluster's hosts. For instructions, see Deploy an OVF or OVA Template in the VMware vSphere online documentation.

To deploy the gateway .ova image

1. Deploy the .ova image to one of the hosts in the cluster.
2. Make sure the data stores that you choose for the root disk and the cache are available to all hosts in the cluster.

(Optional) Add Override Options for Other VMs on Your Cluster

If you have other VMs running on your cluster, you might want to set the cluster values specifically for each VM. For instructions, see Customize an Individual Virtual Machine in the VMware vSphere online documentation.

To add override options for other VMs on your cluster

1. On the **Summary** page in VMware vSphere, choose your cluster to open the cluster page, and then choose **Configure**.
2. Choose the **Configuration** tab, and then choose **VM Overrides**.
3. Add a new VM override option to change each value.

   Set the following values for each option under **vSphere HA - VM Monitoring**:

   - **VM Monitoring**: **Override Enabled - VM and Application Monitoring**
   - **VM monitoring sensitivity**: **Override Enabled - VM and Application Monitoring**
   - **VM Monitoring**: **Custom**
   - **Failure interval**: **30 seconds**
   - **Minimum uptime**: **120 seconds**
   - **Maximum per-VM resets**: **5**
   - **Maximum resets time window**: **Within 1 hrs**

Activate Your Gateway

After the .ova is deployed in your VMware environment, activate your gateway using the Storage Gateway console. For instructions, see Review settings and activate your Amazon FSx File Gateway.

Test Your VMware High Availability Configuration

After you activate your gateway, test your configuration.

To test your VMware HA configuration

1. Open the Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. On the navigation pane, choose **Gateways**, and then choose the gateway that you want to test for VMware HA.
3. For **Actions**, choose **Verify VMware HA**.
4. In the **Verify VMware High Availability Configuration** box that appears, choose **OK**.

> **Note**
>
> Testing your VMware HA configuration reboots your gateway VM and interrupts connectivity to your gateway. The test might take a few minutes to complete.

If the test is successful, the status of **Verified** appears in the details tab of the gateway in the console.

5. Choose **Exit**.

You can find information about VMware HA events in the Amazon CloudWatch log groups. For more information, see Getting FSx File Gateway health logs with CloudWatch log groups.

Getting an activation key for your gateway

To receive an activation key for your gateway, make a web request to the gateway virtual machine (VM). The VM returns a redirect that contains the activation key, which is passed as one of the parameters for the `ActivateGateway` API action to specify the configuration of your gateway. For more information, see ActivateGateway in the *Storage Gateway API Reference*.

> **Note**
>
> Gateway activation keys expire in 30 minutes if unused.

The request that you make to the gateway VM includes the AWS Region where the activation occurs. The URL that's returned by the redirect in the response contains a query string parameter called `activationkey`. This query string parameter is your activation key. The format of the query string looks like the following: `http://`*`gateway_ip_address`*`/?activationRegion=`*`activation_region`*. The output of this query returns both activation region and key.

The URL also includes `vpcEndpoint`, the VPC Endpoint ID for gateways that connect using the VPC endpoint type.

> **Note**
>
> The AWS Storage Gateway Hardware Appliance, VM image templates, and Amazon EC2 Amazon Machine Images (AMI) come preconfigured with the HTTP services necessary to receive and respond to the web requests described on this page. It's not required or recommended to install any additional services on your gateway.

Topics

Linux (curl)

The following examples show you how to get an activation key using Linux (curl).

> **Note**
>
> Replace the highlighted variables with actual values for your gateway. Acceptable values are as follows:
>
> - *`gateway_ip_address`* - The IPv4 address of your gateway, for example `172.31.29.201`
> - *`gateway_type`* - The type of gateway you want to activate, such as `STORED`, `CACHED`, `VTL`, `FILE_S3`, or `FILE_FSX_SMB`.
> - *`region_code`* - The Region where you want to activate your gateway. See Regional endpoints in the *AWS General Reference Guide*. If this parameter is not specified, or if the value provided is misspelled or doesn't match a valid region, the command will default to the `us-east-1` region.
> - *`vpc_endpoint`* - The VPC endpoint name for your gateway, for example `vpce-050f90485f28f2fd0-iep0e8vq.storagegateway.us-west-2.vpce.amazonaws.com`.

To get the activation key for a public endpoint:

```
curl "http://gateway_ip_address/?activationRegion=region_code&no_redirect"
```

To get the activation key for a VPC endpoint:

```
curl "http://gateway_ip_address/?
activationRegion=region_code&vpcEndpoint=vpc_endpoint&no_redirect"
```

Linux (bash/zsh)

The following example shows you how to use Linux (bash/zsh) to fetch the HTTP response, parse HTTP headers, and get the activation key.

```
function get-activation-key() {
  local ip_address=$1
  local activation_region=$2
  if [[ -z "$ip_address" || -z "$activation_region" || -z "$gateway_type" ]]; then
    echo "Usage: get-activation-key ip_address activation_region gateway_type"
    return 1
  fi

  if redirect_url=$(curl -f -s -S -w '%{redirect_url}' "http://$ip_address/?
activationRegion=$activation_region&gatewayType=$gateway_type"); then
    activation_key_param=$(echo "$redirect_url" | grep -oE 'activationKey=[A-Z0-9-]+')
    echo "$activation_key_param" | cut -f2 -d=
  else
    return 1
  fi
}
```

Microsoft Windows PowerShell

The following example shows you how to use Microsoft Windows PowerShell to fetch the HTTP response, parse HTTP headers, and get the activation key.

```
function Get-ActivationKey {
  [CmdletBinding()]
```

```
  Param(
    [parameter(Mandatory=$true)][string]$IpAddress,
    [parameter(Mandatory=$true)][string]$ActivationRegion,
    [parameter(Mandatory=$true)][string]$GatewayType
  )
  PROCESS {
    $request = Invoke-WebRequest -UseBasicParsing -Uri "http://$IpAddress/?
activationRegion=$ActivationRegion&gatewayType=$GatewayType" -MaximumRedirection 0 -
ErrorAction SilentlyContinue
    if ($request) {
      $activationKeyParam = $request.Headers.Location | Select-String -Pattern
 "activationKey=([A-Z0-9-]+)"
      $activationKeyParam.Matches.Value.Split("=")[1]
    }
  }
}
```

Using your local console

The following example shows you how to use your local console to generate and display an activation key.

To get an activation key for your gateway from your local console

1. Log in to your local console as *admin*.
2. After you log in and see the **AWS Appliance Activation - Configuration** main menu, select `0` to choose **Get activation key**.
3. Select **Storage Gateway** for gateway family option.
4. When prompted, enter the AWS Region where you want to activate your gateway.
5. Enter 1 for Public or 2 for VPC endpoint as the network type.
6. Enter 1 for Standard or 2 for Federal Information Processing Standard (FIPS) as the endpoint Type.

Using AWS Direct Connect with Storage Gateway

AWS Direct Connect links your internal network to the Amazon Web Services Cloud. By using AWS Direct Connect with Storage Gateway, you can create a connection for high-throughput workload needs, providing a dedicated network connection between your on-premises gateway and AWS.

Storage Gateway uses public endpoints. With an AWS Direct Connect connection in place, you can create a public virtual interface to allow traffic to be routed to the Storage Gateway endpoints. The public virtual interface bypasses internet service providers in your network path. The Storage Gateway service public endpoint can be in the same AWS Region as the AWS Direct Connect location, or it can be in a different AWS Region.

The following illustration shows an example of how AWS Direct Connect works with Storage Gateway.
network architecture showing Storage Gateway connected to the cloud using AWS direct connect.

The following procedure assumes that you have created a functioning gateway.

To use AWS Direct Connect with Storage Gateway

1. Create and establish an AWS Direct Connect connection between your on-premises data center and your Storage Gateway endpoint. For more information about how to create a connection, see Getting Started with AWS Direct Connect in the *AWS Direct Connect User Guide.*
2. Connect your on-premises Storage Gateway appliance to the AWS Direct Connect router.
3. Create a public virtual interface, and configure your on-premises router accordingly. For more information, see Creating a Virtual Interface in the *AWS Direct Connect User Guide.*

For details about AWS Direct Connect, see What is AWS Direct Connect? in the *AWS Direct Connect User Guide*.

Active Directory service account permission requirements

If you plan to use Microsoft Active directory to provide user authenticated access to the file systems on your AWS Storage Gateway, you need to make sure that you have an Active Directory service account, and that the service account has delegated permissions to join computers to your domain. A service account is an Active Directory user account that has been delegated permission to perform certain tasks. You provide the username and password credentials for this account when you join a Storage Gateway to your Active Directory domain.

The Active Directory service account must be delegated the following permissions in the OU to which you are joining your gateway:

- Ability to create and delete computer objects
- Ability to reset passwords

- Ability to modify permissions
- Ability to restrict accounts from reading and writing data
- Validated ability to read and write Account Restrictions
- Validated ability to write to the service principal name
- Validated ability to write to the DNS host name

These represent the minimum set of permissions that are required to join computer objects to your Active Directory. For more information, see the Microsoft 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.

Getting the IP address for your gateway appliance

After you choose a host and deploy your gateway VM, you connect and activate your gateway. To do this, you need the IP address of your gateway VM. You get the IP address from your gateway's local console. You log in to the local console and get the IP address from the top of the console page.

For gateways deployed on-premises, you can also get the IP address from your hypervisor. For Amazon EC2 gateways, you can also get the IP address of your Amazon EC2 instance from the Amazon EC2 Management Console. To find how to get your gateway's IP address, see one of the following:

- VMware host: Accessing the Gateway Local Console with VMware ESXi
- HyperV host: Access the Gateway Local Console with Microsoft Hyper-V
- Linux Kernel-based Virtual Machine (KVM) host: Accessing the Gateway Local Console with Linux KVM
- EC2 host: Getting an IP Address from an Amazon EC2 Host

When you locate the IP address, take note of it. Then return to the Storage Gateway console and type the IP address into the console.

Getting an IP Address from an Amazon EC2 Host

To get the IP address of the Amazon EC2 instance your gateway is deployed on, log in to the EC2 instance's local console. Then get the IP address from the top of the console page. For instructions, see .

You can also get the IP address from the Amazon EC2 Management Console. We recommend using the public IP address for activation. To get the public IP address, use procedure 1. If you choose to use the elastic IP address instead, see procedure 2.

Procedure 1: To connect to your gateway using the public IP address

1. Open the Amazon EC2 console at https://console.aws.amazon.com/ec2/.
2. In the navigation pane, choose **Instances**, and then select the EC2 instance that your gateway is deployed on.
3. Choose the **Description** tab at the bottom, and then note the public IP. You use this IP address to connect to the gateway. Return to the Storage Gateway console and type in the IP address.

If you want to use the elastic IP address for activation, use the procedure following.

Procedure 2: To connect to your gateway using the elastic IP address

1. Open the Amazon EC2 console at https://console.aws.amazon.com/ec2/.
2. In the navigation pane, choose **Instances**, and then select the EC2 instance that your gateway is deployed on.
3. Choose the **Description** tab at the bottom, and then note the **Elastic IP** value. You use this elastic IP address to connect to the gateway. Return to the Storage Gateway console and type in the elastic IP address.
4. After your gateway is activated, choose the gateway that you just activated, and then choose the **VTL devices** tab in the bottom panel.
5. Get the names of all your VTL devices.
6. For each target, run the following command to configure the target.

   ```
   iscsiadm -m node -o new -T [$TARGET_NAME] -p [$Elastic_IP]:3260
   ```

7. For each target, run the following command to log in.

   ```
   iscsiadm -m node -p [$ELASTIC_IP]:3260 --login
   ```

Your gateway is now connected using the elastic IP address of the EC2 instance.

Understanding Storage Gateway resources and resource IDs

In Storage Gateway, the primary resource is a *gateway* but other resource types is *file share*. File shares are referred to as *subresources* and they don't exist unless they are associated with a gateway.

These resources and subresources have unique Amazon Resource Names (ARNs) associated with them as shown in this table.

Resource Type	ARN Format
Gateway ARN	`arn:aws:storagegateway:` *`region`*`:`*`account-id`*`:gateway/`*`gateway-id`*
File Share ARN	`arn:aws:storagegateway:` *`region`*`:`*`account-id`*`:share/`*`share-id`*

Working with Resource IDs

When you create a resource, Storage Gateway assigns the resource a unique resource ID. This resource ID is part of the resource ARN. A resource ID takes the form of a resource identifier, followed by a hyphen, and a unique combination of eight letters and numbers. For example, a gateway ID is of the form `sgw-12A3456B` where `sgw` is the resource identifier for gateways.

Storage Gateway resource IDs are in uppercase. However, when you use these resource IDs with the Amazon EC2 API, Amazon EC2 expects resource IDs in lowercase. You must change your resource ID to lowercase to use it with the EC2 API. For example, in Storage Gateway the ID for a volume might be `vol-1122AABB`. When you use this ID with the EC2 API, you must change it to `vol-1122aabb`. Otherwise, the EC2 API might not behave as expected.

Important

IDs for Storage Gateway volumes and Amazon EBS snapshots created from gateway volumes are changing to a longer format. Starting in December 2016, all new volumes and

snapshots will be created with a 17-character string. Starting in April 2016, you will be able to use these longer IDs so you can test your systems with the new format. For more information, see Longer EC2 and EBS Resource IDs.
For example, a volume ARN with the longer volume ID format will look like this: `arn:aws:storagegateway:us-west-2:111122223333:gateway/sgw-12A3456B/volume/vol-1122AABBCCDDEEFFG`.
A snapshot ID with the longer ID format will look like this: `snap-78e226633445566ee`.
For more information, see Announcement: Heads-up – Longer Storage Gateway volume and snapshot IDs coming in 2016.

Tagging Storage Gateway resources

In Storage Gateway, you can use tags to manage your resources. Tags let you add metadata to your resources and categorize your resources to make them easier to manage. Each tag consists of a key-value pair, which you define. You can add tags to gateways, volumes, and virtual tapes. You can search and filter these resources based on the tags you add.

As an example, you can use tags to identify Storage Gateway resources used by each department in your organization. You might tag gateways and volumes used by your accounting department like this: (`key=department` and `value=accounting`). You can then filter with this tag to identify all gateways and volumes used by your accounting department and use the information to determine cost. For more information, see Using Cost Allocation Tags and Working with Tag Editor.

If you archive a virtual tape that is tagged, the tape maintains its tags in the archive. Similarly, if you retrieve a tape from the archive to another gateway, the tags are maintained in the new gateway.

For File Gateway, you can use tags to control access to resources. For information about how to do this, see Using tags to control access to your gateway and resources.

Tags don't have any semantic meaning but rather are interpreted as strings of characters.

The following restrictions apply to tags:

- Tag keys and values are case-sensitive.
- The maximum number of tags for each resource is 50.
- Tag keys cannot begin with `aws:`. This prefix is reserved for AWS use.

- Valid characters for the key property are UTF-8 letters and numbers, space, and special characters + - = . _ : / and @.

Working with tags

You can work with tags by using the Storage Gateway console, the Storage Gateway API, or the Storage Gateway Command Line Interface (CLI). The following procedures show you how to add, edit, and delete a tag on the console.

To add a tag

1. Open the Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. In the navigation pane, choose the resource you want to tag.

   For example, to tag a gateway, choose **Gateways**, and then choose the gateway you want to tag from the list of gateways.
3. Choose **Tags**, and then choose **Add/edit tags**.
4. In the **Add/edit tags** dialog box, choose **Create tag**.
5. Type a key for **Key** and a value for **Value**. For example, you can type **`Department`** for the key and **`Accounting`** for the value.

   > **Note**
   >
   > You can leave the **Value** box blank.

6. Choose **Create Tag** to add more tags. You can add multiple tags to a resource.
7. When you're done adding tags, choose **Save**.

To edit a tag

1. Open the Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. Choose the resource whose tag you want to edit.
3. Choose **Tags** to open the **Add/edit tags** dialog box.
4. Choose the pencil icon next to the tag you want to edit, and then edit the tag.

5. When you're done editing the tag, choose **Save**.

To delete a tag

1. Open the Storage Gateway console at https://console.aws.amazon.com/storagegateway/home.
2. Choose the resource whose tag you want to delete.
3. Choose **Tags**, and then choose **Add/edit tags** to open the **Add/edit tags** dialog box.
4. Choose the **X** icon next to the tag you want to delete, and then choose **Save**.

Working with open-source components for AWS Storage Gateway

This section describes the third-party tools and licenses that we depend on to deliver AWS Storage Gateway functionality.

Topics

Open-source components for Storage Gateway

Several third-party tools and licenses are used to deliver functionality for Volume Gateway, Tape Gateway, and Amazon S3 File Gateway.

Use the following links to download source code for certain open-source software components that are included with AWS Storage Gateway software:

- For Storage Gateway appliances deployed on VMware ESXi: sources.tar
- For Storage Gateway appliances deployed on Microsoft Hyper-V: sources_hyperv.tar
- For Storage Gateway appliances deployed on Linux Kernel-based Virtual Machine (KVM): sources_KVM.tar

This product includes software developed by the OpenSSL project for use in the OpenSSL Toolkit (http://www.openssl.org/). For the relevant licenses for all dependent third-party tools, see Third-Party Licenses.

Open-source components for Amazon FSx File Gateway

Several third-party tools and licenses are used to deliver Amazon FSx File Gateway (FSx File Gateway) functionality.

Use the following links to download the source code for certain open-source software components that are included with FSx File Gateway software:

- For Amazon FSx File Gateway 2021-07-07 Release: sgw-file-fsx-smb-open-source.tgz
- For Amazon FSx File Gateway 2021-04-06 Release: sgw-file-fsx-smb-20210406-open-source.tgz

This product includes software developed by the OpenSSL project for use in the OpenSSL Toolkit (http://www.openssl.org/). For the relevant licenses for all dependent third-party tools, see the following links:

- For Amazon FSx File Gateway 2021-07-07 Release: Third-Party License.
- For Amazon FSx File Gateway 2021-04-06 Release: Third-Party License.

Limits and quotas for Amazon FSx File Gateway

Quotas for Amazon FSx file systems

The following table lists minimum and maximum limits and quotas for Amazon FSx file systems.

Resource	Limit per Amazon FSx file system
Maximum number of tags	50 tags
Maximum retention period for automated backups	90 days
Maximum number of backup copy requests in progress to a single destination Region per account.	5 requests

Resource	Limit per Amazon FSx file system
Minimum storage capacity for SSD file systems	32 GiB
Minimum storage capacity for HDD file systems	2,000 GiB
Maximum storage capacity for SSD and HDD file systems	64 TiB
Minimum throughput capacity	8 MBps
Maximum throughput capacity	2,048 MBps
Maximum number of Amazon FSx file shares	100,000

Recommended local disk sizes for your gateway

The following table recommends sizes for local disk storage for each AWS Storage Gateway in your deployment.

Gateway Type	Cache (Minimum)	Cache (Maximum)	Other Required Local Disks
FSx File Gateway	150 GiB	64 TiB	—

Note

You can configure one or more local drives for your cache up to the maximum capacity. When adding cache to an existing FSx File Gateway, it is important to create new disks on your virtual host (hypervisor or Amazon EC2 instance). Do not change the size of existing disks if the disks have been previously allocated as a cache.

API Reference for Storage Gateway

In addition to using the console, you can use the AWS Storage Gateway API to programmatically configure and manage your gateways. This section describes the AWS Storage Gateway operations, request signing for authentication and the error handling. For information about the regions and endpoints available for Storage Gateway, see AWS Storage Gateway Endpoints and Quotas in the *AWS General Reference*.

> **Note**
>
> You can also use the AWS SDKs when developing applications with Storage Gateway. The AWS SDKs for Java, .NET, and PHP wrap the underlying Storage Gateway API, simplifying your programming tasks. For information about downloading the SDK libraries, see Sample Code Libraries.

Topics

AWS Storage Gateway Required Request Headers

This section describes the required headers that you must send with every POST request to AWS Storage Gateway. You include HTTP headers to identify key information about the request including the operation you want to invoke, the date of the request, and information that indicates the authorization of you as the sender of the request. Headers are case insensitive and the order of the headers is not important.

The following example shows headers that are used in the ActivateGateway operation.

```
POST / HTTP/1.1
Host: storagegateway.us-east-2.amazonaws.com
Content-Type: application/x-amz-json-1.1
```

```
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120425/us-east-2/
storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target,
 Signature=9cd5a3584d1d67d57e61f120f35102d6b3649066abdd4bf4bbcf05bd9f2f8fe2
x-amz-date: 20120912T120000Z
x-amz-target: StorageGateway_20120630.ActivateGateway
```

The following are the headers that must include with your POST requests to AWS Storage Gateway. Headers shown below that begin with "x-amz" are AWS-specific headers. All other headers listed are common header used in HTTP transactions.

Header	Description
`Authorization`	The authorization header contains several of pieces of information about the request that allow AWS Storage Gateway to determine if the request is a valid action for the requester. The format of this header is as follows (line breaks added for readability): `Authorization: AWS4-HMAC_SHA456 Credentials= YourAccessKey /yyymmdd/region/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature= CalculatedSignature` In the preceding syntax, you specify *YourAccessKey*, the year, month, and day (*yyyymmdd*), the *region*, and the *CalculatedSignature*. The format of the authorization header is dictated by the requirements of the AWS V4 Signing process. The details of signing are discussed in the topic Signing Requests.
`Content-Type`	Use `application/x-amz-json-1.1` as the content type for all requests to AWS Storage Gateway. `Content-Type: application/x-amz-json-1.1`
`Host`	Use the host header to specify the AWS Storage Gateway endpoint where you send your request. For example, `storagegateway.us-`

Header	Description
	`east-2.amazonaws.com` is the endpoint for the US East (Ohio) region. For more information about the endpoints available for AWS Storage Gateway, see AWS Storage Gateway Endpoints and Quotas in the *AWS General Reference*. `Host: storagegateway. region.amazonaws.com`
`x-amz-date`	You must provide the time stamp in either the HTTP `Date` header or the AWS `x-amz-date` header. (Some HTTP client libraries don't let you set the `Date` header.) When an `x-amz-date` header is present, the AWS Storage Gateway ignores any `Date` header during the request authentication. The `x-amz-date` format must be ISO8601 Basic in the YYYYMMDD'T'HHMMSS'Z' format. If both the `Date` and `x-amz-date` header are used, the format of the Date header does not have to be ISO8601. `x-amz-date: YYYYMMDD'T'HHMMSS'Z'`
`x-amz-target`	This header specifies the version of the API and the operation that you are requesting. The target header values are formed by concatenating the API version with the API name and are in the following format. `x-amz-target: StorageGateway_ APIversion .operationName` The *operationName* value (e.g. "ActivateGateway") can be found from the API list, API Reference for Storage Gateway.

Signing Requests

Storage Gateway requires that you authenticate every request you send by signing the request. To sign a request, you calculate a digital signature using a cryptographic hash function. A cryptographic hash is a function that returns a unique hash value based on the input. The input to

the hash function includes the text of your request and your secret access key. The hash function returns a hash value that you include in the request as your signature. The signature is part of the `Authorization` header of your request.

After receiving your request, Storage Gateway recalculates the signature using the same hash function and input that you used to sign the request. If the resulting signature matches the signature in the request, Storage Gateway processes the request. Otherwise, the request is rejected.

Storage Gateway supports authentication using AWS Signature Version 4. The process for calculating a signature can be broken into three tasks:

- Task 1: Create a Canonical Request

  Rearrange your HTTP request into a canonical format. Using a canonical form is necessary because Storage Gateway uses the same canonical form when it recalculates a signature to compare with the one you sent.
- Task 2: Create a String to Sign

  Create a string that you will use as one of the input values to your cryptographic hash function. The string, called the *string to sign*, is a concatenation of the name of the hash algorithm, the request date, a *credential scope* string, and the canonicalized request from the previous task. The *credential scope* string itself is a concatenation of date, region, and service information.
- Task 3: Create a Signature

  Create a signature for your request by using a cryptographic hash function that accepts two input strings: your *string to sign* and a *derived key*. The *derived key* is calculated by starting with your secret access key and using the *credential scope* string to create a series of Hash-based Message Authentication Codes (HMACs).

Example Signature Calculation

The following example walks you through the details of creating a signature for ListGateways. The example could be used as a reference to check your signature calculation method.

The example assumes the following:

- The time stamp of the request is "Mon, 10 Sep 2012 00:00:00" GMT.
- The endpoint is the US East (Ohio) region.

The general request syntax (including the JSON body) is:

```
POST / HTTP/1.1
Host: storagegateway.us-east-2.amazonaws.com
x-amz-Date: 20120910T000000Z
Authorization: SignatureToBeCalculated
Content-type: application/x-amz-json-1.1
x-amz-target: StorageGateway_20120630.ListGateways
{}
```

The canonical form of the request calculated for Task 1: Create a Canonical Request is:

```
POST
/

content-type:application/x-amz-json-1.1
host:storagegateway.us-east-2.amazonaws.com
x-amz-date:20120910T000000Z
x-amz-target:StorageGateway_20120630.ListGateways

content-type;host;x-amz-date;x-amz-target
44136fa355b3678a1146ad16f7e8649e94fb4fc21fe77e8310c060f61caaff8a
```

The last line of the canonical request is the hash of the request body. Also, note the empty third line in the canonical request. This is because there are no query parameters for this API (or any Storage Gateway APIs).

The *string to sign* for Task 2: Create a String to Sign is:

```
AWS4-HMAC-SHA256
20120910T000000Z
20120910/us-east-2/storagegateway/aws4_request
92c0effa6f9224ac752ca179a04cecbede3038b0959666a8160ab452c9e51b3e
```

The first line of the *string to sign* is the algorithm, the second line is the time stamp, the third line is the *credential scope*, and the last line is a hash of the canonical request from Task 1.

For Task 3: Create a Signature, the *derived key* can be represented as:

```
derived key = HMAC(HMAC(HMAC(HMAC("AWS4" + YourSecretAccessKey,"20120910"),"us-
east-2"),"storagegateway"),"aws4_request")
```

If the secret access key, wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY, is used, then the calculated signature is:

```
6d4c40b8f2257534dbdca9f326f147a0a7a419b63aff349d9d9c737c9a0f4c81
```

The final step is to construct the `Authorization` header. For the demonstration access key AKIAIOSFODNN7EXAMPLE, the header (with line breaks added for readability) is:

```
Authorization: AWS4-HMAC-SHA256 Credential=AKIAIOSFODNN7EXAMPLE/20120910/us-east-2/
storagegateway/aws4_request,
SignedHeaders=content-type;host;x-amz-date;x-amz-target,
Signature=6d4c40b8f2257534dbdca9f326f147a0a7a419b63aff349d9d9c737c9a0f4c81
```

Error Responses

Topics

- Exceptions
- Operation Error Codes
- Error Responses

This section provides reference information about AWS Storage Gateway errors. These errors are represented by an error exception and an operation error code. For example, the error exception `InvalidSignatureException` is returned by any API response if there is a problem with the request signature. However, the operation error code `ActivationKeyInvalid` is returned only for the ActivateGateway API.

Depending on the type of error, Storage Gateway may return only just an exception, or it may return both an exception and an operation error code. Examples of error responses are shown in the Error Responses.

Exceptions

The following table lists AWS Storage Gateway API exceptions. When an AWS Storage Gateway operation returns an error response, the response body contains one of these exceptions. The `InternalServerError` and `InvalidGatewayRequestException` return one of the operation error codes Operation Error Codes message codes that give the specific operation error code.

Exception	Message	HTTP Status Code
`IncompleteSignatureException`	The specified signature is incomplete.	400 Bad Request
`InternalFailure`	The request processing has failed due to some unknown error, exception or failure.	500 Internal Server Error
`InternalServerError`	One of the operation error code messages Operation Error Codes.	500 Internal Server Error
`InvalidAction`	The requested action or operation is invalid.	400 Bad Request
`InvalidClientTokenId`	The X.509 certificate or AWS Access Key ID provided does not exist in our records.	403 Forbidden
`InvalidGatewayRequestException`	One of the operation error code messages in Operation Error Codes.	400 Bad Request
`InvalidSignatureException`	The request signature we calculated does not match the signature you provided. Check your AWS Access Key and signing method.	400 Bad Request
`MissingAction`	The request is missing an action or operation parameter.	400 Bad Request

Exception	Message	HTTP Status Code
MissingAuthenticationToken	The request must contain either a valid (registered) AWS Access Key ID or X.509 certificate.	403 Forbidden
RequestExpired	The request is past the expiration date or the request date (either with 15 minute padding), or the request date occurs more than 15 minutes in the future.	400 Bad Request
SerializationException	An error occurred during serialization. Check that your JSON payload is well-formed.	400 Bad Request
ServiceUnavailable	The request has failed due to a temporary failure of the server.	503 Service Unavailable
SubscriptionRequiredException	The AWS Access Key Id needs a subscription for the service.	400 Bad Request
ThrottlingException	Rate exceeded.	400 Bad Request
TooManyRequests	Too many requests.	429 Too Many Requests
UnknownOperationException	An unknown operation was specified. Valid operations are listed in Storage Gateway API Actions.	400 Bad Request
UnrecognizedClientException	The security token included in the request is invalid.	400 Bad Request
ValidationException	The value of an input parameter is bad or out of range.	400 Bad Request

Operation Error Codes

The following table shows the mapping between AWS Storage Gateway operation error codes and APIs that can return the codes. All operation error codes are returned with one of two general exceptions—`InternalServerError` and `InvalidGatewayRequestException`—described in Exceptions.

Operation Error Code	Message	Operations That Return this Error Code
`ActivationKeyExpired`	The specified activation key has expired.	ActivateGateway
`ActivationKeyInvalid`	The specified activation key is invalid.	ActivateGateway
`ActivationKeyNotFound`	The specified activation key was not found.	ActivateGateway
`BandwidthThrottleScheduleNotFound`	The specified bandwidth throttle was not found.	DeleteBandwidthRateLimit
`CannotExportSnapshot`	The specified snapshot cannot be exported.	CreateCachediSCSIVolume CreateStorediSCSIVolume
`InitiatorNotFound`	The specified initiator was not found.	DeleteChapCredentials
`DiskAlreadyAllocated`	The specified disk is already allocated.	AddCache AddUploadBuffer AddWorkingStorage CreateStorediSCSIVolume
`DiskDoesNotExist`	The specified disk does not exist.	AddCache AddUploadBuffer

Operation Error Code	Message	Operations That Return this Error Code
		AddWorkingStorage CreateStorediSCSIVolume
DiskSizeNotGigAligned	The specified disk is not gigabyte-aligned.	CreateStorediSCSIVolume
DiskSizeGreaterThanVolumeMaxSize	The specified disk size is greater than the maximum volume size.	CreateStorediSCSIVolume
DiskSizeLessThanVolumeSize	The specified disk size is less than the volume size.	CreateStorediSCSIVolume
DuplicateCertificateInfo	The specified certificate information is a duplicate.	ActivateGateway
FileSystemAssociationEndpointConfigurationConflict	Existing File System Association endpoint configuration conflicts with specified configuration.	AssociateFileSystem
FileSystemAssociationEndpointIpAddressAlreadyInUse	The specified endpoint IP address is already in use.	AssociateFileSystem
FileSystemAssociationEndpointIpAddressMissing	File System Association Endpoint IP address is missing.	AssociateFileSystem

Operation Error Code	Message	Operations That Return this Error Code
FileSystemAssociationNotFound	The specified file system association was not found.	UpdateFileSystemAssociation DisassociateFileSystem DescribeFileSystemAssociations
FileSystemNotFound	The specified file system was not found.	AssociateFileSystem

Operation Error Code	Message	Operations That Return this Error Code
GatewayInternalError	A gateway internal error occurred.	AddCache AddUploadBuffer AddWorkingStorage CreateCachediSCSIVolume CreateSnapshot CreateStorediSCSIVolume CreateSnapshotFromVolumeRecoveryPoint DeleteBandwidthRateLimit DeleteChapCredentials DeleteVolume DescribeBandwidthRateLimit DescribeCache DescribeCachediSCSIVolumes DescribeChapCredentials DescribeGatewayInformation DescribeMaintenanceStartTime DescribeSnapshotSchedule DescribeStorediSCSIVolumes DescribeWorkingStorage ListLocalDisks

Operation Error Code	Message	Operations That Return this Error Code
		ListVolumes ListVolumeRecoveryPoints ShutdownGateway StartGateway UpdateBandwidthRateLimit UpdateChapCredentials UpdateMaintenanceStartTime UpdateGatewaySoftwareNow UpdateSnapshotSchedule

Operation Error Code	Message	Operations That Return this Error Code
GatewayNotConnected	The specified gateway is not connected.	AddCache AddUploadBuffer AddWorkingStorage CreateCachediSCSIVolume CreateSnapshot CreateStorediSCSIVolume CreateSnapshotFromVolumeRecoveryPoint DeleteBandwidthRateLimit DeleteChapCredentials DeleteVolume DescribeBandwidthRateLimit DescribeCache DescribeCachediSCSIVolumes DescribeChapCredentials DescribeGatewayInformation DescribeMaintenanceStartTime DescribeSnapshotSchedule DescribeStorediSCSIVolumes DescribeWorkingStorage ListLocalDisks

Operation Error Code	Message	Operations That Return this Error Code
		ListVolumes ListVolumeRecoveryPoints ShutdownGateway StartGateway UpdateBandwidthRateLimit UpdateChapCredentials UpdateMaintenanceStartTime UpdateGatewaySoftwareNow UpdateSnapshotSchedule

Operation Error Code	Message	Operations That Return this Error Code
GatewayNotFound	The specified gateway was not found.	AddCache AddUploadBuffer AddWorkingStorage CreateCachediSCSIVolume CreateSnapshot CreateSnapshotFromVolumeRecoveryPoint CreateStorediSCSIVolume DeleteBandwidthRateLimit DeleteChapCredentials DeleteGateway DeleteVolume DescribeBandwidthRateLimit DescribeCache DescribeCachediSCSIVolumes DescribeChapCredentials DescribeGatewayInformation DescribeMaintenanceStartTime DescribeSnapshotSchedule DescribeStorediSCSIVolumes DescribeWorkingStorage

Operation Error Code	Message	Operations That Return this Error Code
		ListLocalDisks ListVolumes ListVolumeRecoveryPoints ShutdownGateway StartGateway UpdateBandwidthRateLimit UpdateChapCredentials UpdateMaintenanceStartTime UpdateGatewaySoftwareNow UpdateSnapshotSchedule

Operation Error Code	Message	Operations That Return this Error Code
`GatewayProxyNetworkConnectionBusy`	The specified gateway proxy network connection is busy.	AddCache AddUploadBuffer AddWorkingStorage CreateCachediSCSIVolume CreateSnapshot CreateSnapshotFromVolumeRecoveryPoint CreateStorediSCSIVolume DeleteBandwidthRateLimit DeleteChapCredentials DeleteVolume DescribeBandwidthRateLimit DescribeCache DescribeCachediSCSIVolumes DescribeChapCredentials DescribeGatewayInformation DescribeMaintenanceStartTime DescribeSnapshotSchedule DescribeStorediSCSIVolumes DescribeWorkingStorage ListLocalDisks

Operation Error Code	Message	Operations That Return this Error Code
		ListVolumes ListVolumeRecoveryPoints ShutdownGateway StartGateway UpdateBandwidthRateLimit UpdateChapCredentials UpdateMaintenanceStartTime UpdateGatewaySoftwareNow UpdateSnapshotSchedule

Operation Error Code	Message	Operations That Return this Error Code
`InternalError`	An internal error occurred.	ActivateGateway AddCache AddUploadBuffer AddWorkingStorage CreateCachediSCSIVolume CreateSnapshot CreateSnapshotFromVolumeRecoveryPoint CreateStorediSCSIVolume DeleteBandwidthRateLimit DeleteChapCredentials DeleteGateway DeleteVolume DescribeBandwidthRateLimit DescribeCache DescribeCachediSCSIVolumes DescribeChapCredentials DescribeGatewayInformation DescribeMaintenanceStartTime DescribeSnapshotSchedule DescribeStorediSCSIVolumes

Operation Error Code	Message	Operations That Return this Error Code
		DescribeWorkingStorage ListLocalDisks ListGateways ListVolumes ListVolumeRecoveryPoints ShutdownGateway StartGateway UpdateBandwidthRateLimit UpdateChapCredentials UpdateMaintenanceStartTime UpdateGatewayInformation UpdateGatewaySoftwareNow UpdateSnapshotSchedule

Operation Error Code	Message	Operations That Return this Error Code
`InvalidParameters`	The specified request contains invalid parameters.	ActivateGateway AddCache AddUploadBuffer AddWorkingStorage CreateCachediSCSIVolume CreateSnapshot CreateSnapshotFromVolumeRecoveryPoint CreateStorediSCSIVolume DeleteBandwidthRateLimit DeleteChapCredentials DeleteGateway DeleteVolume DescribeBandwidthRateLimit DescribeCache DescribeCachediSCSIVolumes DescribeChapCredentials DescribeGatewayInformation DescribeMaintenanceStartTime DescribeSnapshotSchedule DescribeStorediSCSIVolumes

Operation Error Code	Message	Operations That Return this Error Code
		DescribeWorkingStorage ListLocalDisks ListGateways ListVolumes ListVolumeRecoveryPoints ShutdownGateway StartGateway UpdateBandwidthRateLimit UpdateChapCredentials UpdateMaintenanceStartTime UpdateGatewayInformation UpdateGatewaySoftwareNow UpdateSnapshotSchedule
`LocalStorageLimitExceeded`	The local storage limit was exceeded.	AddCache AddUploadBuffer AddWorkingStorage
`LunInvalid`	The specified LUN is invalid.	CreateStorediSCSIVolume

Operation Error Code	Message	Operations That Return this Error Code
MaximumVolumeCountExceeded	The maximum volume count was exceeded.	CreateCachediSCSIVolume CreateStorediSCSIVolume DescribeCachediSCSIVolumes DescribeStorediSCSIVolumes
NetworkConfigurationChanged	The gateway network configuration has changed.	CreateCachediSCSIVolume CreateStorediSCSIVolume

Operation Error Code	Message	Operations That Return this Error Code
`NotSupported`	The specified operation is not supported.	ActivateGateway AddCache AddUploadBuffer AddWorkingStorage CreateCachediSCSIVolume CreateSnapshot CreateSnapshotFromVolumeRecoveryPoint CreateStorediSCSIVolume DeleteBandwidthRateLimit DeleteChapCredentials DeleteGateway DeleteVolume DescribeBandwidthRateLimit DescribeCache DescribeCachediSCSIVolumes DescribeChapCredentials DescribeGatewayInformation DescribeMaintenanceStartTime DescribeSnapshotSchedule DescribeStorediSCSIVolumes

Operation Error Code	Message	Operations That Return this Error Code
		DescribeWorkingStorage ListLocalDisks ListGateways ListVolumes ListVolumeRecoveryPoints ShutdownGateway StartGateway UpdateBandwidthRateLimit UpdateChapCredentials UpdateMaintenanceStartTime UpdateGatewayInformation UpdateGatewaySoftwareNow UpdateSnapshotSchedule
`OutdatedGateway`	The specified gateway is out of date.	ActivateGateway
`SnapshotInProgressException`	The specified snapshot is in progress.	DeleteVolume
`SnapshotIdInvalid`	The specified snapshot is invalid.	CreateCachediSCSIVolume CreateStorediSCSIVolume
`StagingAreaFull`	The staging area is full.	CreateCachediSCSIVolume CreateStorediSCSIVolume

Operation Error Code	Message	Operations That Return this Error Code
`TargetAlreadyExists`	The specified target already exists.	CreateCachediSCSIVolume CreateStorediSCSIVolume
`TargetInvalid`	The specified target is invalid.	CreateCachediSCSIVolume CreateStorediSCSIVolume DeleteChapCredentials DescribeChapCredentials UpdateChapCredentials
`TargetNotFound`	The specified target was not found.	CreateCachediSCSIVolume CreateStorediSCSIVolume DeleteChapCredentials DescribeChapCredentials DeleteVolume UpdateChapCredentials

Operation Error Code	Message	Operations That Return this Error Code
UnsupportedOperationForGatewayType	The specified operation is not valid for the type of the gateway.	AddCache AddWorkingStorage CreateCachediSCSIVolume CreateSnapshotFromVolumeRecoveryPoint CreateStorediSCSIVolume DeleteSnapshotSchedule DescribeCache DescribeCachediSCSIVolumes DescribeStorediSCSIVolumes DescribeUploadBuffer DescribeWorkingStorage ListVolumeRecoveryPoints
VolumeAlreadyExists	The specified volume already exists.	CreateCachediSCSIVolume CreateStorediSCSIVolume
VolumeIdInvalid	The specified volume is invalid.	DeleteVolume
VolumeInUse	The specified volume is already in use.	DeleteVolume

Operation Error Code	Message	Operations That Return this Error Code
VolumeNotFound	The specified volume was not found.	CreateSnapshot CreateSnapshotFromVolumeRecoveryPoint DeleteVolume DescribeCachediSCSIVolumes DescribeSnapshotSchedule DescribeStorediSCSIVolumes UpdateSnapshotSchedule
VolumeNotReady	The specified volume is not ready.	CreateSnapshot CreateSnapshotFromVolumeRecoveryPoint

Error Responses

When there is an error, the response header information contains:

- Content-Type: application/x-amz-json-1.1
- An appropriate 4xx or 5xx HTTP status code

The body of an error response contains information about the error that occurred. The following sample error response shows the output syntax of response elements common to all error responses.

```
{
    "__type": "String",
    "message": "String",
    "error":
        { "errorCode": "String",
```

```
          "errorDetails": "String"
        }
}
```

The following table explains the JSON error response fields shown in the preceding syntax.

__type

One of the exceptions from Exceptions.

Type: String

error

Contains API-specific error details. In general errors (i.e., not specific to any API), this error information is not shown.

Type: Collection

errorCode

One of the operation error codes .

Type: String

errorDetails

This field is not used in the current version of the API.

Type: String

message

One of the operation error code messages.

Type: String

Error Response Examples

The following JSON body is returned if you use the DescribeStorediSCSIVolumes API and specify a gateway ARN request input that does not exist.

```
{
  "__type": "InvalidGatewayRequestException",
```

```
  "message": "The specified volume was not found.",
  "error": {
    "errorCode": "VolumeNotFound"
  }
}
```

The following JSON body is returned if Storage Gateway calculates a signature that does not match the signature sent with a request.

```
{
  "__type": "InvalidSignatureException",
  "message": "The request signature we calculated does not match the signature you
 provided."
}
```

Storage Gateway API Actions

For a list of Storage Gateway operations, see Actions in the *AWS Storage Gateway API Reference*.

Document history for the Amazon FSx File Gateway User Guide

- **API version**: 2013-06-30
- **Latest documentation update**: June 06, 2024

The following table describes important changes in each release of this user guide after April 2018. For notification about updates to this documentation, you can subscribe to an RSS feed.

Change	Description	Date
Notice of availability change for FSx File Gateway	Amazon FSx File Gateway is no longer available to new customers. Existing customers of FSx File Gateway can continue to use the service normally. For capabilities similar to FSx File Gateway, visit this blog post.	October 28, 2024
Notice of availability change for FSx File Gateway	AWS Storage Gateway's FSx File Gateway will no longer be available to new customers starting 10/28/24. To use the service, you must sign up prior to that date. Existing customers of FSx File Gateway can continue to use the service normally. For capabilities similar to FSx File Gateway, visit this blog post.	September 26, 2024
Added option to turn maintenance updates on or off	Storage Gateway receives regular maintenance updates that can include operating	June 6, 2024

	system and software upgrades, fixes to address stability, performance, and security, and access to new features. You can now configure a setting to turn these updates on or off for each individual gateway in your deployment. For more information, see Managing gateway updates using the AWS Storage Gateway console.	
Updated recommended CloudWatch alarms	The CloudWatch `HealthNotifications` alarm now applies to and is recommended for all gateway types and host platforms. Recommended configuration settings have also been updated for `HealthNotifications` and `AvailabilityNotifications`. For more information see Understanding CloudWatch alarms.	October 2, 2023

Added GatewayClockOutOfSync troubleshooting tips	The *Troubleshooting: File Gateway issues section* now includes troubleshooting guidelines to help diagnose problems you may encounter if your gateway system clock is not synchronized with the AWS Storage Gateway server time. For more information, see Error: GatewayClockOutOfSync.	October 19, 2022
Added Active Directory Join Domain troubleshooting tips	The *Troubleshooting: File Gateway issues section* now includes troubleshooting guidelines to help diagnose problems you may encounter when trying to join your gateway to an Active Directory domain. For more information, see Troubleshooting: Active Directory domain issues.	October 19, 2022
Updated gateway creation procedures	The procedure for creating a new gateway has been updated to reflect changes in the Storage Gateway console. For more information, see Create and activate an Amazon S3 File Gateway.	October 12, 2021

Multiple file system support	Amazon FSx File Gateway now supports up to five attached Amazon FSx file systems. For more information, see Attach an Amazon FSx for Windows File Server file system.	July 7, 2021
Amazon FSx soft storage quota support	Amazon FSx File Gateway now supports soft storage quotas (which warn you when users surpass their data limits) when writing to attached Amazon FSx file systems where storage quotas are configured. Hard quotas (which enforce data limits by denying write access) are not supported. Soft quotas work for all users except the Amazon FSx admin user. For more information about setting up storage quotas, see Storage quotas in the *Amazon FSx for Windows File Server User Guide*.	July 7, 2021

New guide	In addition to the original File Gateway (now known as Amazon S3 File Gateway), Storage Gateway provides Amazon FSx File Gateway (FSx File Gateway). FSx File Gateway provides low latency and efficient access to in-cloud FSx for Windows File Server file shares from your on-premises facility. For more information, see What is Amazon FSx File Gateway?	April 27, 2021
FedRAMP compliance	Storage Gateway is now FedRAMP compliant. For more information, see Compliance validation for Storage Gateway.	November 24, 2020
File Gateway migration	File Gateway now provides a documented process for replacing an existing File Gateway with a new File Gateway. For more informati on, see Replacing a File Gateway with a new File Gateway.	October 30, 2020
File Gateway cold cache read performance 4x increase	Storage Gateway has increased cold cache read performance 4x. For more information, see Performance guidance for File Gateways.	August 31, 2020

Order the hardware appliance through the console	You can now order the hardware appliance through the AWS Storage Gateway console. For more informati on, see Using the AWS Storage Gateway Hardware Appliance.	August 12, 2020
Support for Federal Informati on Processing Standard (FIPS) endpoints in new AWS Regions	You can now activate a gateway with FIPS endpoints in the US East (Ohio), US E ast (N. Virginia), US West (N. California), US West (Oregon), and Canada (Central) Regions. For more informati on, see AWS Storage Gateway endpoints and quotas in the *AWS General Reference*.	July 31, 2020
File Gateway local cache storage 4x increase	Storage Gateway now supports a local cache of up to 64 TB for File Gateway, improving performance for on-premises applications by providing low-latency access to larger working datasets. For more information, see Recommended local disk sizes for your gateway in the *Storage Gateway User Guide*.	July 7, 2020
View Amazon CloudWatc h alarms in the Storage Gateway console	You can now view CloudWatc h alarms in the Storage Gateway console. For more information, see Understan ding CloudWatch alarms.	May 29, 2020

Support for Federal Information Processing Standard (FIPS) endpoints	You can now activate a gateway with FIPS endpoints in the AWS GovCloud (US) Regions. To choose a FIPS endpoint for a File Gateway, see Choosing a service endpoint.	May 22, 2020
New AWS Regions	Storage Gateway is now available in the Africa (Cape Town) and Europe (Milan) Regions. For more information, see AWS Storage Gateway endpoints and quotas in the *AWS General Reference*.	May 7, 2020
Support for S3 Intelligent-Tiering storage class	Storage Gateway now supports S3 Intelligent-Tiering storage class. The S3 Intelligent-Tiering storage class optimizes storage costs by automatically moving data to the most cost-effective storage access tier, without performance impact or operational overhead. For more information, see Storage class for automatically optimizing frequently and infrequently accessed objects in the *Amazon Simple Storage Service User Guide*.	April 30, 2020

New AWS Region	Storage Gateway is now available in the AWS GovCloud (US-East) Region. For more information, see AWS Storage Gateway Endpoints and Quotas in the *AWS General Reference*.	March 12, 2020
Support for Linux Kernel-based Virtual Machine (KVM) hypervisor	Storage Gateway now provides the ability to deploy an on-premises gateway on the KVM virtualization platform. Gateways deployed on KVM have all the same functionality and features as the existing on-premises gateways. For more information, see Supported Hypervisors and Host Requirements in the *Storage Gateway User Guide*.	February 4, 2020

Support for VMware vSphere High Availability	Storage Gateway now provides support for high availability on VMware to help protect storage workloads against hardware, hypervisor, or network failures. For more information, see Using VMware vSphere High Availability with Storage Gateway in the *Storage Gateway User Guide*. This release also includes performance improvements. For more information, see Performance in the *Storage Gateway User Guide*.	November 20, 2019
Support for Amazon CloudWatch Logs	You can now configure File Gateways with Amazon CloudWatch Log Groups to get notified about errors and the health of your gateway and its resources. For more information, see Getting Notified About Gateway Health and Errors With Amazon CloudWatch Log Groups in the *Storage Gateway User Guide.*	September 4, 2019

New AWS Region	Storage Gateway is now available in the Asia Pacific (Hong Kong) Region. For more information, see AWS Storage Gateway Endpoints and Quotas in the *AWS General Reference*.	August 14, 2019
New AWS Region	Storage Gateway is now available in the Middle East (Bahrain) Region. For more information, see AWS Storage Gateway Endpoints and Quotas in the *AWS General Reference*.	July 29, 2019
Support for activating a gateway in a virtual private cloud (VPC)	You can now activate a gateway in a VPC. You can create a private connection between your on-premises software appliance and cloud-based storage infrastructure. For more information, see Activating a Gateway in a Virtual Private Cloud.	June 20, 2019
File Gateway support for tag-based authorization	File Gateway now supports tag-based authorization. You can control access to File Gateway resources based on the tags on those resources. You can also control access based on the tags that can be passed in an IAM request condition. For more information, see Controlling Access to File Gateway Resources.	March 4, 2019

Availability of AWS Storage Gateway Hardware Appliance in Europe	The AWS Storage Gateway Hardware Appliance is now available in Europe. For more information, see AWS Storage Gateway Hardware Appliance Regions in the *AWS General Reference*. In addition, you can now increase the useable storage on the AWS Storage Gateway Hardware Appliance from 5 TB to 12 TB and replace the installed copper network card with a 10-gigabi t fiber optic network card. For more information, see Setting Up Your Hardware Appliance.	February 25, 2019
Support for AWS Storage Gateway Hardware Appliance	The AWS Storage Gateway Hardware Appliance includes Storage Gateway software preinstalled on a third-party server. You can manage the appliance from the AWS Management Console. The appliance can host file, tape, and Volume Gateways. For more information, see Using the Storage Gateway Hard ware Appliance.	September 18, 2018

Earlier updates

The following table describes important changes in each release of the *AWS Storage Gateway User Guide* before May 2018.

Change	Description	Date Changed
New AWS Region	Tape Gateway is now available in the Asia Pacific (Singapore) Region. For detailed information, see AWS Regions that support Storage Gateway.	April 3, 2018
New AWS Region	Storage Gateway is now available in the Europe (Paris) Region. For detailed information, see AWS Regions that support Storage Gateway.	December 18, 2017
Support for VMware ESXi Hypervisor version 6.5	AWS Storage Gateway now supports VMware ESXi Hypervisor version 6.5. This is in addition to version 4.1, 5.0, 5.1, 5.5, and 6.0. For more information, see Supported hypervisors and host requirements.	September 13, 2017
File Gateway support for Microsoft Hyper-V hypervisor	You can now deploy a File Gateway on a Microsoft Hyper-V hypervisor. For information, see Supported hypervisors and host requirements.	June 22, 2017
New AWS Region	Storage Gateway is now available in the Asia Pacific (Mumbai) Region. For detailed information, see AWS Regions that support Storage Gateway.	May 02, 2017
Support for File Gateways on Amazon EC2	AWS Storage Gateway now provides the ability to deploy a File Gateway in Amazon EC2. You can launch a File Gateway in Amazon EC2 using the Storage Gateway Amazon Machine Image (AMI) now available as a community AMI. For information about how to create a File Gateway and deploy it on an EC2 instance, see Create and activate an Amazon FSx File Gateway. For information about how to launch a File Gateway AMI, see Deploy a default Amazon EC2 host for FSx File Gateway. In addition, File Gateway now supports HTTP proxy configuration. For more information, see Routing	February 08, 2017

Change	Description	Date Changed
	your gateway deployed on Amazon EC2 through an HTTP proxy.	
New AWS Region	Storage Gateway is now available in the Europe (London) Region. For detailed information, see AWS Regions that support Storage Gateway.	December 13, 2016
New AWS Region	Storage Gateway is now available in the Canada (Central) Region. For detailed information, see AWS Regions that support Storage Gateway.	December 08, 2016
Support for File Gateway	In addition to Volume Gateways and Tape Gateway, Storage Gateway now provides File Gateway. File Gateway combines a service and virtual software appliance, allowing you to store and retrieve objects in Amazon S3 using industry-standard file protocols such as Network File System (NFS). The gateway provides access to objects in Amazon S3 as files on an NFS mount point.	November 29, 2016