



Amazon S3 File Gateway User Guide

AWS Storage Gateway



API Version 2013-06-30

Copyright © 2025 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

AWS Storage Gateway: Amazon S3 File Gateway User Guide

Copyright © 2025 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

What is Amazon S3 File Gateway	1
How S3 File Gateway works	2
Getting Started with AWS Storage Gateway	5
Sign up for Amazon Web Services	5
Create an IAM user with administrator privileges	6
Accessing AWS Storage Gateway	7
AWS Regions that support Storage Gateway	8
File Gateway setup requirements	9
Prerequisites	9
Hardware and storage requirements	10
Hardware requirements for on-premises VMs	10
Requirements for Amazon EC2 instance types	10
Storage requirements	11
Network and firewall requirements	12
Port requirements	13
Networking and firewall requirements for the hardware appliance	28
Allowing gateway access through firewall and routers	31
Configuring security group	35
Supported hypervisors and host requirements	35
Supported NFS and SMB clients for File Gateway	36
Supported file system operations	37
Managing local disks	37
Deciding the amount of local disk storage	38
Add cache storage	39
Using ephemeral storage with EC2 gateways	40
Using the hardware appliance	42
Setting up your hardware appliance	43
Physically installing your hardware appliance	44
Accessing the hardware appliance console	46
Configuring hardware appliance network parameters	48
Activating your hardware appliance	49
Creating a gateway on your hardware appliance	51
Configuring a gateway IP address on the hardware appliance	52
Removing gateway software from your hardware appliance	54

Deleting your hardware appliance	55
Creating your gateway	57
Overview - Gateway Activation	57
Set up gateway	57
Connect to AWS	57
Review and activate	58
Overview - Gateway Configuration	58
Overview - Storage Resources	58
Create an S3 File Gateway	58
Set up an Amazon S3 File Gateway	59
Connect your Amazon S3 File Gateway to AWS	60
Review settings and activate your Amazon S3 File Gateway	61
Configure your Amazon S3 File Gateway	62
Activating a gateway in a VPC	64
Create a VPC endpoint for Storage Gateway	65
Creating a file share	67
Avoid unanticipated costs	68
Encrypt objects stored by File Gateway	69
Create an NFS file share	70
Create NFS file share with default configuration	71
Create NFS file share with custom configuration	75
Create an SMB file share	82
Create SMB file share with default configuration	83
Create SMB file share with custom configuration	89
Mounting and using your file share	100
Mount your NFS file share on your client	100
Mount your SMB file share on your client	102
Using file shares on buckets with pre-existing objects	105
Test your S3 File Gateway	106
Managing your Amazon S3 File Gateway	108
Edit basic gateway information	109
Granting access and permissions	110
Granting access to an S3 bucket	110
Cross-service confused deputy prevention	114
Using a file share for cross-account access	115
Delete a file share	116

Editing gateway SMB settings	118
Set gateway security level	119
Configure Active Directory authentication	120
Provide guest access	123
Configure local groups	123
Set file share visibility	124
Edit SMB file share settings	125
Limit SMB file share access	126
Change file share encryption method	127
Edit NFS file share settings	129
Edit NFS file share metadata defaults	131
Limit NFS file share access	132
Refreshing Amazon S3 bucket object cache	133
Configure an automated cache refresh schedule using the Storage Gateway console	134
Configure an automated cache refresh schedule using AWS Lambda with an Amazon CloudWatch rule	135
Perform a manual cache refresh using the Storage Gateway console	138
Perform a manual cache refresh using the Storage Gateway API	139
Using S3 Object Lock	140
File share status	140
Gateway status	141
Managing bandwidth	142
Edit bandwidth-rate-limit schedule	143
Using the AWS SDK for Java	145
Using the AWS SDK for .NET	147
Using the AWS Tools for Windows PowerShell	149
Monitoring Storage Gateway	151
Understanding CloudWatch alarms	151
Create recommended CloudWatch alarms	153
Create a custom CloudWatch alarm	154
Monitoring your S3 File Gateway	156
Getting S3 File Gateway health logs	156
Using Amazon CloudWatch metrics	158
Getting notified about file operations	160
Understanding gateway metrics	168
Understanding file share metrics	174

Understanding S3 File Gateway audit logs	177
Create a cache report	182
Manage cache reports	185
Understanding cache reports	187
Maintaining your gateway	189
Managing gateway updates	189
Update frequency and expected behavior	190
Turn maintenance updates on or off	191
Modify the gateway maintenance window schedule	192
Apply an update manually	193
Performing maintenance tasks using the local console	194
Accessing the gateway local console	194
Performing tasks on the virtual machine local console	197
Performing tasks on the EC2 local console	212
Shutting down your gateway VM	219
Replacing your existing S3 File Gateway with a new instance	220
Method 1: Migrate cache disk and Gateway ID to replacement instance	221
Method 2: Replacement instance with empty cache disk and new Gateway ID	224
Deleting your gateway and removing resources	226
Deleting Your Gateway by Using the Storage Gateway Console	226
Performance and optimization	228
Basic performance guidance for S3 File Gateway	228
S3 File Gateway performance on Linux clients	229
File Gateway performance on Windows clients	231
Performance guidance for gateways with multiple file shares	232
Maximizing S3 File Gateway throughput	234
Deploy your gateway in the same location as your clients	235
Reduce bottlenecks caused by slow disks	235
Adjust virtual machine resource allocation for CPU, RAM, and cache disks	236
Adjust the SMB security level	238
Use multiple threads and clients to parallelize write operations	238
Turn off automated cache refresh	240
Increase the number of Amazon S3 uploader threads	241
Increase SMB timeout settings	242
Turn on opportunistic locking for compatible applications	242
Adjust gateway capacity according to the size of the working file set	242

Deploy multiple gateways for larger workloads	243
Optimizing S3 File Gateway for SQL Server database backups	244
Deploy your gateway in the same location as your SQL Servers	245
Reduce bottlenecks caused by slow disks	245
Adjust S3 File Gateway virtual machine resource allocation for CPU, RAM, and cache disks	246
Improve SMB client throughput by adjusting the security level of your S3 File Gateway	247
Improve SMB client throughput by splitting SQL backups into multiple files	248
Prevent large file copy failures by increasing SMB timeout settings	249
Increase the number of Amazon S3 uploader threads	249
Turn off automated cache refresh	250
Deploy multiple gateways to support the workload	251
Additional resources for database backup workloads	251
Security	252
Data protection	252
Data encryption	253
Identity and access management	254
Audience	254
Authenticating with identities	255
Managing access using policies	258
How AWS Storage Gateway works with IAM	261
Identity-based policy examples	267
Troubleshooting	270
Using tags to control access to resources	272
Using ACLs for SMB file share access	275
Compliance validation	278
Resilience	279
Infrastructure security	280
AWS Security Best Practices	281
Logging and monitoring	281
Storage Gateway information in CloudTrail	281
Understanding Storage Gateway log file entries	282
Troubleshooting	285
Troubleshooting: gateway offline issues	286
Check the associated firewall or proxy	286
Check for an ongoing SSL or deep-packet inspection of your gateway's traffic	286

Check the IOWaitPercent metric after a reboot or software update	286
Check for a power outage or hardware failure on the hypervisor host	287
Check for issues with an associated cache disk	287
Troubleshooting: Active Directory issues	287
Confirm that the gateway can reach the domain controller by running an nping test	288
Check the DHCP options set for the VPC of your Amazon EC2 gateway instance	289
Confirm that the gateway can resolve the domain by running a dig query	289
Check the domain controller settings and roles	290
Check that the gateway is joined to the nearest domain controller	290
Confirm that Active Directory creates new computer objects in the default organizational unit (OU)	291
Check your domain controller event logs	291
Troubleshooting: gateway activation issues	291
Resolve errors when activating your gateway using a public endpoint	292
Resolve errors when activating your gateway using an Amazon VPC endpoint	295
Resolve errors when activating your gateway using a public endpoint and there is a Storage Gateway VPC endpoint in the same VPC	299
Troubleshooting: on-premises gateway issues	300
Troubleshooting: Open NFS ports	303
Turning on Support access to help troubleshoot your gateway	304
Troubleshooting: Microsoft Hyper-V setup issues	305
Troubleshooting: Amazon EC2 gateway issues	308
Gateway activation hasn't occurred after a few moments	309
Can't find the EC2 gateway instance in the instance list	309
Connect to your Amazon EC2 gateway using the serial console	310
Turning on Support access to help troubleshoot the gateway	310
Troubleshooting: hardware appliance issues	312
How to determine service IP address	312
How to perform a factory reset	312
How to perform a remote restart	313
How to obtain Dell iDRAC support	313
How to find the hardware appliance serial number	313
How to get hardware appliance support	313
Troubleshooting: File Gateway issues	314
Error: 1344 (0x00000540)	315
Error: GatewayClockOutOfSync	315

Error: InaccessibleStorageClass	315
Error: InvalidObjectState	316
Error: ObjectMissing	317
Error: RoleTrustRelationshipInvalid	317
Error: S3AccessDenied	318
Error: DroppedNotifications	318
Notification: HardReboot	319
Notification: Reboot	319
Troubleshooting: Open NFS ports	303
Troubleshooting with CloudWatch metrics	320
Troubleshooting: file share issues	323
File share is stuck in CREATING status	323
Can't create a file share	324
SMB file shares don't allow multiple different access methods	324
Multiple file shares can't write to the mapped S3 bucket	325
Notification for deleted log group when using audit logs	325
Can't upload files into S3 bucket	325
Can't change default encryption to SSE-KMS	325
Changes made directly in an S3 bucket with object versioning turned on may affect what you see in your file share	326
When writing to an S3 bucket with versioning turned on, the Amazon S3 File Gateway may create multiple versions of Amazon S3 objects	327
Changes to an S3 bucket are not reflected in Storage Gateway	328
ACL permissions aren't working as expected	329
Gateway performance declined after a recursive operation	329
High Availability Health Notifications	329
Troubleshooting: high availability issues	330
Health notifications	330
Metrics	331
Best practices	332
Recovering your data	332
Recovering from an unexpected VM shutdown	333
Recovering data from a malfunctioning cache disk	333
Recovering data from an inaccessible data center	333
Managing multipart uploads	334
Unzipping compressed files	334

Copying data from Windows Server	335
Cache disk sizing	335
Multiple file shares and buckets	335
Clean up unnecessary resources	337
Additional resources	338
Host setup	339
Deploy a default Amazon EC2 host for File Gateway	339
Deploy a customized Amazon EC2 host for File Gateway	342
Modify Amazon EC2 instance metadata options	345
Synchronize VM time with Hyper-V or Linux KVM host time	346
Synchronize VM time with VMware host time	347
Configuring network adapters for your gateway	348
Using Storage Gateway with VMware HA	351
Getting activation key	355
Linux (curl)	356
Linux (bash/zsh)	357
Microsoft Windows PowerShell	358
Using your local console	358
File attribute support	359
Using AWS Direct Connect	360
Active Directory permissions	361
Getting the gateway IP address	361
Getting an IP Address from an Amazon EC2 Host	362
Understanding resources and resource IDs	363
Working with Resource IDs	363
Tagging your resources	364
Working with tags	365
Open-source components	366
Open-source components for Storage Gateway	366
Open-source components for Amazon S3 File Gateway	367
Quotas	367
Quotas for file shares	367
Recommended local disk sizes for your gateway	370
Using storage classes	371
Using storage classes with a File Gateway	371
Using the GLACIER storage class with File Gateway	376

Using Kubernetes CSI drivers	376
Working with SMB CSI drivers	377
Working with NFS CSI drivers	381
Terraform module	385
API Reference	387
Required Request Headers	387
Signing Requests	389
Example Signature Calculation	390
Error Responses	392
Exceptions	393
Operation Error Codes	395
Error Responses	415
Actions	417
Document history	418
Earlier updates	434
Release notes	437

What is Amazon S3 File Gateway

Amazon S3 File Gateway – Amazon S3 File Gateway supports a file interface into [Amazon Simple Storage Service \(Amazon S3\)](#) and combines a service and a virtual software appliance. By using this combination, you can store and retrieve objects in Amazon S3 using industry-standard file protocols such as Network File System (NFS) and Server Message Block (SMB). 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. The gateway provides access to objects in S3 as files or file share mount points. With a S3 File Gateway, you can do the following:

- You can store and retrieve files directly using the NFS version 3 or 4.1 protocol.
 - You can store and retrieve files directly using the SMB file system version, 2 and 3 protocol.
 - You can access your data directly in Amazon S3 from any AWS Cloud application or service.
 - You can manage your S3 data using lifecycle policies, cross-region replication, and versioning.
- You can think of a S3 File Gateway as a file system mount on Amazon S3.

A S3 File Gateway simplifies file storage in Amazon S3, integrates to existing applications through industry-standard file system protocols, and provides a cost-effective alternative to on-premises storage. It also provides low-latency access to data through transparent local caching. A S3 File Gateway manages data transfer to and from AWS, buffers applications from network congestion, optimizes and streams data in parallel, and manages bandwidth consumption.

S3 File Gateway integrates with other AWS services, for example:

- Common access management using AWS Identity and Access Management (IAM)
- Encryption using AWS Key Management Service (AWS KMS)
- Monitoring using Amazon CloudWatch (CloudWatch)
- Audit using AWS CloudTrail (CloudTrail)
- Operations using the AWS Management Console and AWS Command Line Interface (AWS CLI)
- Billing and cost management

In the following documentation, you can find a Getting Started section that covers setup information common to all gateways and also gateway-specific setup sections. The Getting Started

section shows you how to deploy, activate, and configure storage for a gateway. The management section shows you how to manage your gateway and resources:

- provides instructions on how to create and use a S3 File Gateway. It shows you how to create a file share, map your drive to an Amazon S3 bucket, and upload files and folders to Amazon S3.
- describes how to perform management tasks for all gateway types and resources.

In this guide, you can primarily find how to work with gateway operations by using the AWS Management Console. If you want to perform these operations programmatically, see the [AWS Storage Gateway API Reference](#).

How Amazon S3 File Gateway works

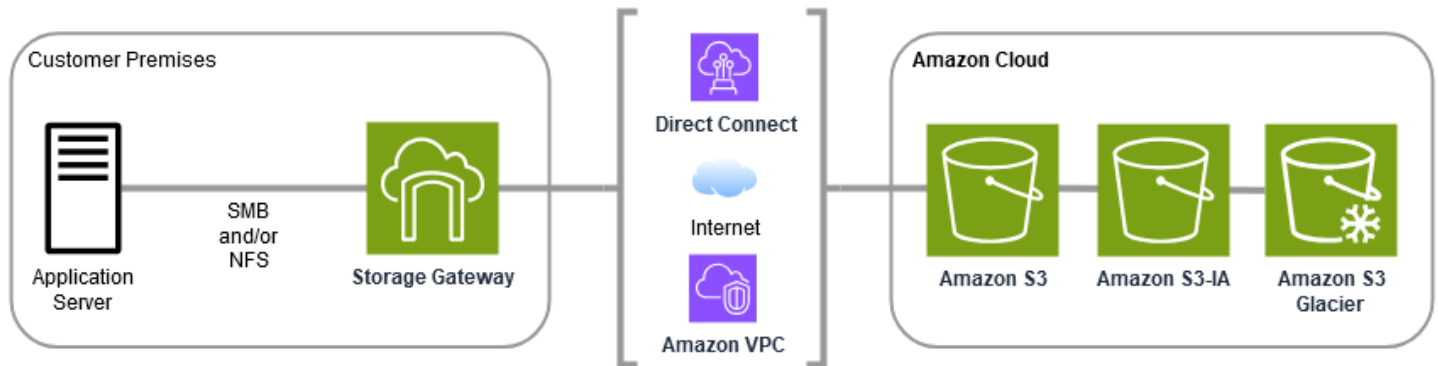
To use an S3 File Gateway, you start by downloading a VM image for the gateway. You then activate the gateway from the AWS Management Console or through the Storage Gateway API. You can also create an S3 File Gateway using an Amazon EC2 image.

After the S3 File Gateway is activated, you create and configure your file share and associate that share with your Amazon Simple Storage Service (Amazon S3) bucket. Doing this makes the share accessible by clients using either the Network File System (NFS) or Server Message Block (SMB) protocol. Files written to a file share become objects in Amazon S3, with the path as the key. There is a one-to-one mapping between files and objects, and the gateway asynchronously updates the objects in Amazon S3 as you change the files. Existing objects in the Amazon S3 bucket appear as files in the file system, and the key becomes the path. Objects are encrypted with Amazon S3–server-side encryption keys (SSE-S3). All data transfer is done through HTTPS.

When sending HTTPS data upload requests to the Amazon S3, File Gateway populates the Content-MD5 header with the MD5 checksum of the data being uploaded. The use of this header causes Amazon S3 to return a failure in case of mismatch between the MD5 checksum computed by Amazon S3 and the value received from the File Gateway. If such a failure is returned, the File Gateway resends the request.

The service optimizes data transfer between the gateway and AWS using multipart parallel uploads or byte-range downloads, to better use the available bandwidth. Local cache is maintained to provide low latency access to the recently accessed data and reduce data egress charges. CloudWatch metrics provide insight into resource use on the VM and data transfer to and from AWS. CloudTrail tracks all API calls.

With S3 File Gateway storage, you can do such tasks as ingesting cloud workloads to Amazon S3, performing backups and archiving, tiering, and migrating storage data to the AWS Cloud. The following diagram provides an overview of file storage deployment for Storage Gateway.



S3 File Gateway converts files to S3 objects when uploading files to Amazon S3. The interaction between file operations performed against files shares on S3 File Gateway and S3 objects requires certain operations to be carefully considered when converting between files and objects.

Common file operations change file metadata, which results in the deletion of the current S3 object and the creation of a new S3 object. The following table shows example file operations and the impact on S3 objects.

File operation	S3 object impact	Storage class implication
Rename file	Replaces existing S3 object and creates a new S3 object for each file	Early deletion fees and retrieval fees may apply
Rename folder	Replaces all existing S3 objects and creates new S3 objects for each folder and files in the folder structure	Early deletion fees and retrieval fees may apply
Change file/folder permissions	Replaces existing S3 object and creates a new S3 object for each file or folder	Early deletion fees and retrieval fees may apply

File operation	S3 object impact	Storage class implication
Change file/folder ownership	Replaces existing S3 object and creates a new S3 object for each file or folder	Early deletion fees and retrieval fees may apply
Append to a file	Replaces existing S3 object and creates a new S3 object for each file	Early deletion fees and retrieval fees may apply

When a file is written to the S3 File Gateway by an NFS or SMB client, the File Gateway uploads the file's data to Amazon S3 followed by its metadata, (ownerships, timestamps, etc.). Uploading the file data creates an S3 object, and uploading the metadata for the file updates the metadata for the S3 object. This process creates another version of the object, resulting in two versions of an object. If S3 Versioning is turned on, both versions will be stored.

When a file is modified in the S3 File Gateway by an NFS or SMB client after it has been uploaded to Amazon S3, the S3 File Gateway uploads the new or modified data instead of uploading the whole file. The file modification results in a new version of the S3 object being created.

When the S3 File Gateway uploads larger files, it might need to upload smaller chunks of the file before the client is done writing to the S3 File Gateway. Some reasons for this include freeing up cache space or a high rate of writes to a file share. This can result in multiple versions of an object in the S3 bucket.

You should monitor your S3 bucket to determine how many versions of an object exist before setting up lifecycle policies to move objects to different storage classes. You should configure lifecycle expiration for previous versions to minimize the number of versions you have for an object in your S3 bucket. The use of Same-Region replication (SRR) or Cross-Region replication (CRR) between S3 buckets will increase the storage used.

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 begin 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 <i>AWS IAM Identity Center User Guide</i> .	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 <i>IAM User Guide</i> .		<i>Command Line Interface User Guide</i> .
In IAM (Not recommended)	Use long-term credentials to access AWS.	Following the instructions in Create an IAM user for emergency access in the <i>IAM User Guide</i> .	Configure programmatic access by Manage access keys for IAM users in the <i>IAM User Guide</i> .

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, creating, managing, and deleting file shares, 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.

- 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](#)
- [Hardware and storage requirements](#)
- [Network and firewall requirements](#)
- [Supported hypervisors and host requirements](#)
- [Supported NFS and SMB clients for File Gateway](#)
- [Supported file system operations for File Gateway](#)
- [Managing local disks for your gateway](#)

Prerequisites

Before you set up your Amazon S3 File Gateway (S3 File Gateway), you must meet the following prerequisites:

- 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.

For information about best practices for File Gateway performance, see [Basic performance guidance for S3 File Gateway](#).

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

For more information, see [Maximizing S3 File Gateway throughput](#). For information about how your hardware affects the performance of the gateway VM, see [Quotas for file shares](#).

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*. Application writes are stored in the cache synchronously, and then asynchronously uploaded to durable storage in Amazon S3. If the ephemeral storage is lost because an instance stops before the upload is complete, the data that still resides in the cache and has not yet written to Amazon Simple Storage Service (Amazon S3) can be lost. Before you stop the instance that hosts the gateway, make sure that the `CachePercentDirty` CloudWatch metric is 0. For information about ephemeral storage, see [Using ephemeral storage with EC2 gateways](#). For information about monitoring metrics for your Storage Gateway, see [Monitoring your S3 File Gateway](#).

If you have more than 5 million objects in your S3 bucket and you are using a **gp2** EBS volume, a minimum root EBS volume of 350 GiB is required for acceptable performance of your gateway during start up. Newly-created Amazon EC2 File Gateway instances use **gp3** root volumes by default, which do not have this requirement. For information about how to increase the volume size, see [Modifying an EBS volume using elastic volumes \(console\)](#).

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.

For information about gateway quotas, see [Quotas for file shares](#).

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](#)
- [Networking and firewall requirements for the Storage Gateway Hardware Appliance](#)
- [Allowing AWS Storage Gateway access through firewalls and routers](#)
- [Configuring security groups for your Amazon EC2 gateway instance](#)

Port requirements

S3 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 NFS or SMB clients, VPC endpoints, or Microsoft Active Directory.

For S3 File Gateway, you only need to use Microsoft Active Directory when you want 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 S3 File Gateway

Network Element	From	To	Protocol	Port	Inbound	Outbound	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

Network Element	From	To	Protocol	Port	Inbound	Outbound	Required	Notes
								during activation of a Storage Gateway appliance . A Storage Gateway VM doesn't require port 80 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

Network Element	From	To	Protocol	Port	Inbound	Outbound	Required	Notes
								Gateway Management Console, the host from which you connect to the console must have access to your gateway's port 80.
Web browser	Storage Gateway VM	AWS	TCP HTTPS	443	✓	✓	✓	AWS Management Console (all other operations)

Network Element	From	To	Protocol	Port	Inbound	Outbound	Required	Notes
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	Outbound	Required	Notes
NTP	Storage Gateway VM	Network Time Protocol (NTP) server	TCP & UDP NTP	123	✓	✓	✓	<p>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:</p> <ul style="list-style-type: none"> • 0.amazon.pool.ntp.org • 1.amazon.pool.ntp.org • 2.amazon.pool.ntp.org

Network Element	From	To	Protocol	Port	Inbound	Outbound	Required	Notes
								<div><div><div>• 3.amazon.pool.ntp.org</div></div><div><div><div><div><div><div></div><div></div></div><div><div><div>Note</div><div>Not required for gateways hosted on Amazon EC2.</div></div></div></div></div></div></div></div>

Network Element	From	To	Protocol	Port	Inbound	Outbound	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	Outbound	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	Outbound	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	Outbound	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	Outbound	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. *Required for SMB only.
Microsoft Active Directory	Storage Gateway VM	Active Directory server	UDP NetBIOS	137	✓	✓	✓*	Name service *Required for SMBv1 only.
Microsoft Active Directory	Storage Gateway VM	Active Directory server	UDP NetBIOS	138	✓	✓	✓*	Datagram service *Required for SMBv1 only.

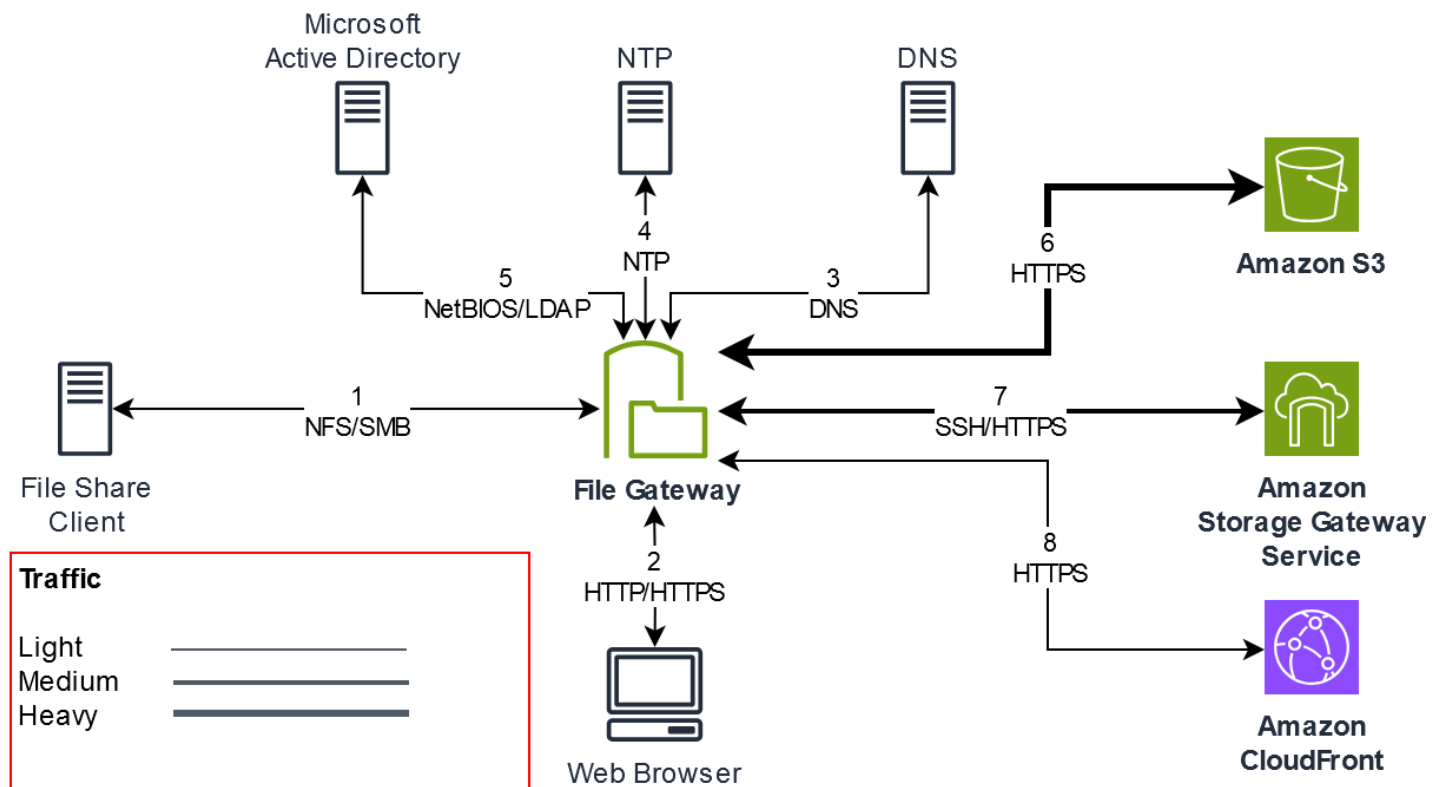
Network Element	From	To	Protocol	Port	Inbound	Outbound	Required	Notes
Microsoft Active Directory	Storage Gateway VM	Active Directory server	TCP & UDP LDAP	389	✓	✓	✓*	Directory System Agent (DSA) client connection *Required for SMB only.
Microsoft Active Directory	Storage Gateway VM	Active Directory server	TCP & UDP Kerberos	88	✓	✓	✓*	Kerberos *Required for SMB only.
Microsoft Active Directory	Storage Gateway VM	Active Directory server	TCP Distributed Computing Environment/End Point Mapper (DCE/EMAP)	135	✓	✓	✓*	RPC *Required for SMB only.

Network Element	From	To	Protocol	Port	Inbound	Outbound	Required	Notes
File share client	NFS Client	Storage Gateway VM	TCP or UDP Data NFSv3	111	✓	✓	✓*	File sharing data transfer (for NFS v3 only) *Required for NFS only.
File share client	NFS Client	Storage Gateway VM	TCP or UDP NFS	2049	✓	✓	✓*	File sharing data transfer *Required for NFS v3 & v4 only.
File share client	NFS Client	Storage Gateway VM	TCP or UDP NFSv3	20048	✓	✓	✓*	File sharing data transfer *Required for NFSv3 only

Network Element	From	To	Protocol	Port	Inbound	Outbound	Required	Notes
File share client	NFS Client	Storage Gateway VM	TCP or UDP NFSv3	8750	✓	✓	✓*	File share quota *Required for NFSv3 only
File share client	SMB Client	Storage Gateway VM	TCP or UDP SMBv2	139	✓	✓	✓*	File sharing data transfer session service *Required for SMB only

Network Element	From	To	Protocol	Port	Inbound	Outbound	Required	Notes
Amazon S3	Storage Gateway VM	Amazon S3 service endpoints	TCP HTTPS	443	✓	✓	✓	For communication from the Storage Gateway VM to the AWS service endpoint. For information about service endpoints, see Allowing AWS Storage Gateway access through firewalls and routers.

The following illustration shows network traffic flow for a basic S3 File Gateway deployment.



Networking and firewall requirements for the Storage Gateway Hardware Appliance

Each Storage Gateway Hardware Appliance requires the following network services:

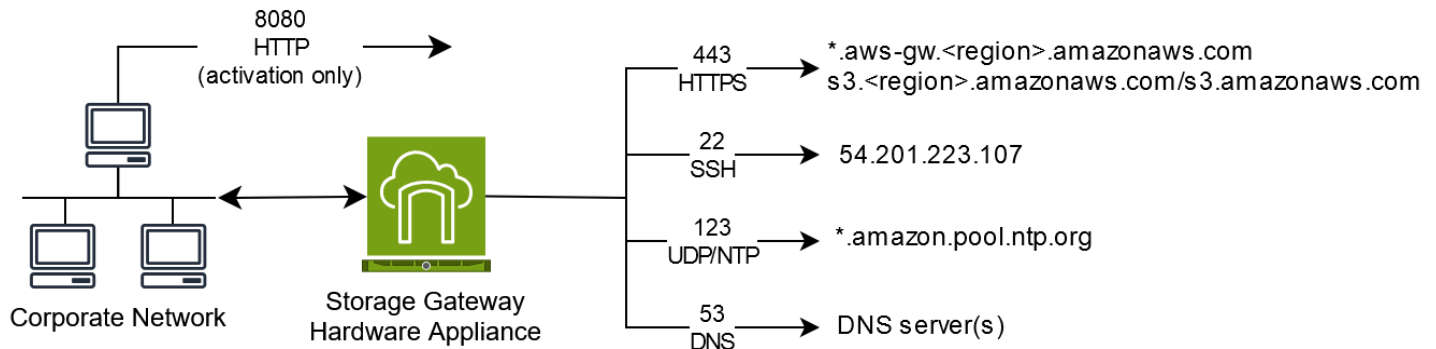
- **Internet access** – an always-on network connection to the internet through any network interface on the server.
- **DNS services** – DNS services for communication between the hardware appliance and DNS server.
- **Time synchronization** – an automatically configured Amazon NTP time service must 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
HTTP	8080	Inbound	AWS	Hardware appliance	Activation (only briefly)

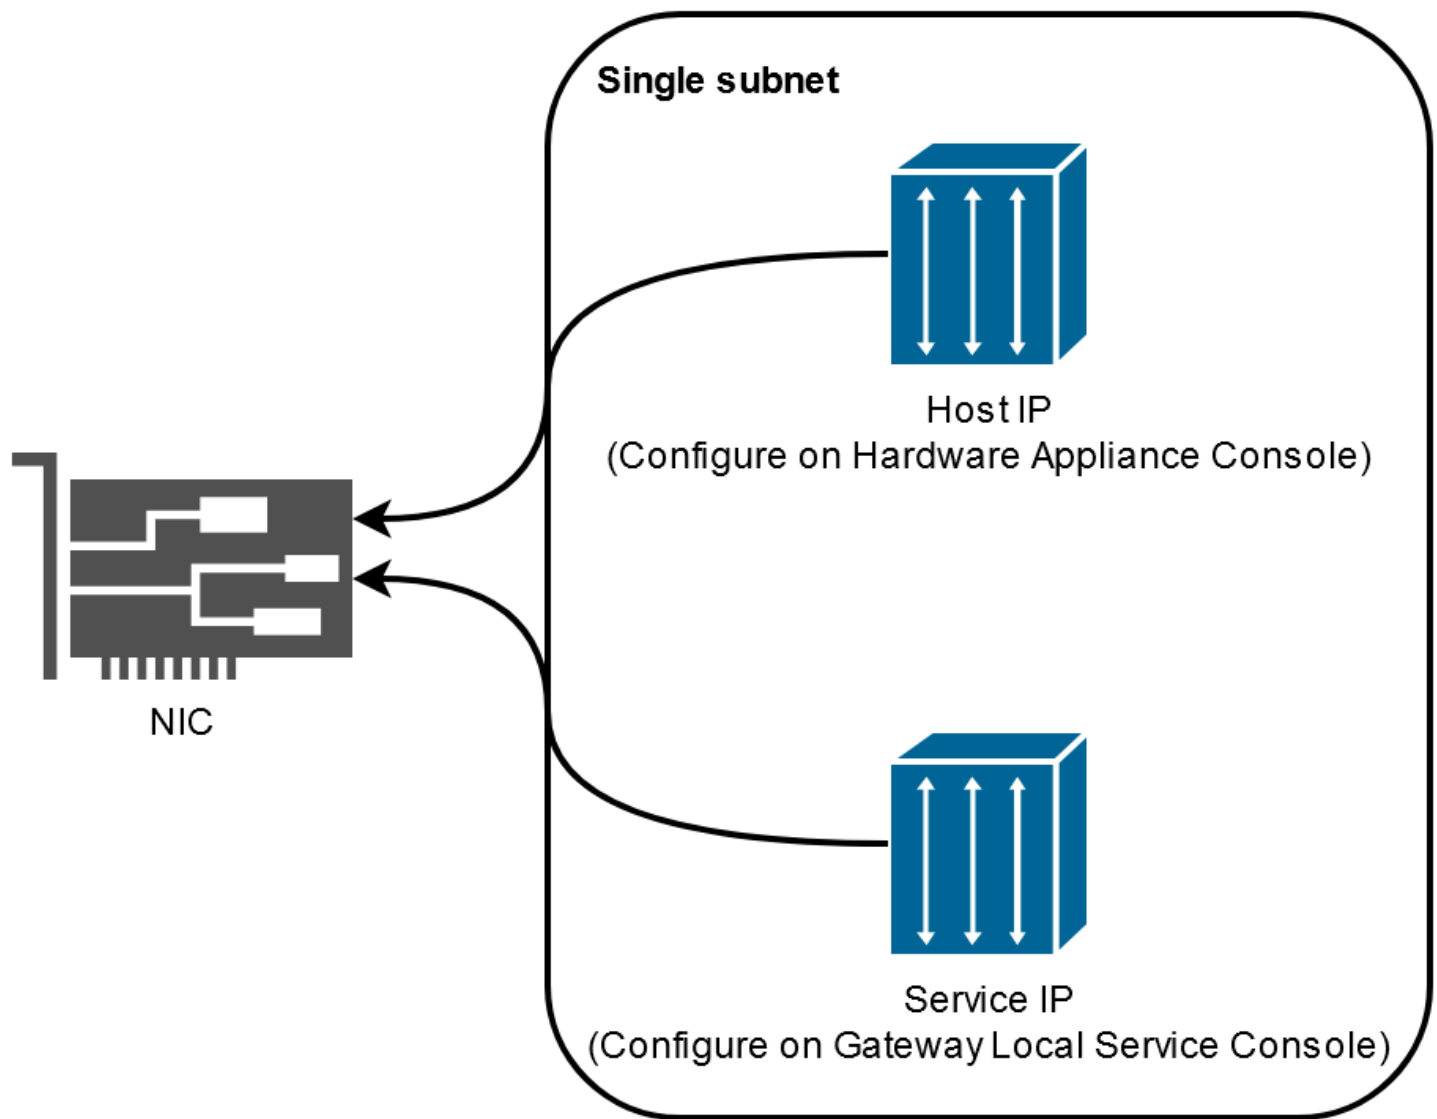
To perform as designed, a hardware appliance requires network and firewall settings as follows:

- 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
```

Amazon S3 service endpoints

Amazon S3 File Gateway requires the following 3 types of endpoints to connect to the Amazon S3 service:

Amazon S3 service endpoint

Note

For this endpoint only, *do not* replace `s3` with `s3-fips` for FIPS-compliant deployments.

```
s3.amazonaws.com
```

Amazon S3 regional endpoints

```
s3.region.amazonaws.com
```

The following example shows an Amazon S3 regional endpoint in the US East (Ohio) Region (`us-east-2`).

```
s3.us-east-2.amazonaws.com
```

The following example shows a FIPS-compliant Amazon S3 regional endpoint in the US West (N. California) Region (`us-west-1`).

```
s3-fips.us-west-1.amazonaws.com
```

The following examples show the Amazon S3 regional endpoints used by AWS GovCloud (US) Regions:

```
s3-fips.us-gov-west-1.amazonaws.com (AWS GovCloud (US-West) Region (FIPS))  
s3-fips.us-gov-east-1.amazonaws.com (AWS GovCloud (US-East) Region (FIPS))  
s3.us-gov-west-1.amazonaws.com (AWS GovCloud (US-West) Region (Standard))
```

```
s3.us-gov-east-1.amazonaws.com (AWS GovCloud (US-East) Region (Standard))
```

Note

If your gateway can't determine the AWS Region where your Amazon S3 bucket is located, this service endpoint defaults to `s3.us-east-1.amazonaws.com`. We recommend that you allow access to the US East (N. Virginia) Region (`us-east-1`) in addition to the AWS Regions where your gateway is activated, and where your Amazon S3 bucket is located.

Amazon S3 bucket endpoints

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

The following example shows an Amazon S3 bucket endpoint for a bucket named `amzn-s3-demo-bucket` in the US East (Ohio) Region (`us-east-2`).

```
amzn-s3-demo-bucket.s3.us-east-2.amazonaws.com
```

The following example shows a FIPS-compliant Amazon S3 bucket endpoint for a bucket named `amzn-s3-demo-bucket1` in the AWS GovCloud (US-East) Region (`us-gov-east-1`).

```
amzn-s3-demo-bucket1.s3-fips.us-gov-east-1.amazonaws.com
```

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](#).

For information about the ports to open for your gateway, see [Port requirements](#).

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 S3 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 NFS and SMB clients for File Gateway

File Gateway supports the following clients:

Operating System Version	Kernel Version	Supported Protocols
Amazon Linux 2023	6.1 LTS	NFSv4.1, NFSv3
Amazon Linux 2	5.10 LTS	NFSv4.1, NFSv3
RHEL 9	5.14	NFSv4.1, NFSv3
RHEL 8.10	4.18	NFSv4.1, NFSv3
SUSE 15	6.4	NFSv4.1, NFSv3
Ubuntu 24.04 LTS	6.8 LTS	NFSv4.1, NFSv3
Ubuntu 22.04 LTS	5.15 LTS	NFSv4.1, NFSv3
Microsoft Windows Server 2025		SMBv2, SMBv3, NFSv3

Operating System Version	Kernel Version	Supported Protocols
Microsoft Windows Server 2022		SMBv2, SMBv3, NFSv3
Microsoft Windows 11		SMBv2, SMBv3, NFSv3
Microsoft Windows 10		SMBv2, SMBv3, NFSv3

Note

Server Message Block (SMB) encryption requires clients that support SMB v3 dialects.

Supported file system operations for File Gateway

Your NFS or SMB client can write, read, delete, and truncate files. When clients send writes to AWS Storage Gateway, it writes to local cache synchronously. Then it writes to Amazon S3 asynchronously through optimized transfers. Reads are first served through the local cache. If data is not available, it's fetched through S3 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 objects from Amazon S3. Directories are managed as folder objects in S3, using the same syntax as in the Amazon S3 console.

HTTP operations such as GET, PUT, UPDATE, and DELETE can modify files in a file share. These operations conform to the atomic create, read, update, and delete (CRUD) functions.

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 S3. 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 S3 File Gateway, consider how much cache disk to allocate. S3 File Gateway uses a least recently used algorithm to automatically evict data from the cache. The cache on an S3 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 S3 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 S3 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

This section describes steps you need to take to prevent data loss when you select an ephemeral disk as storage for your gateway's cache.

Ephemeral disks provide temporary block-level storage for your Amazon EC2 instance. Ephemeral disks are ideal for temporary storage of data that changes frequently, such as data in a gateway's cache storage. 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 Amazon S3. 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 Amazon S3 can be lost. You can prevent such data loss by following the steps before you restart or stop the EC2 instance that hosts your gateway.

 **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.

These steps in this following procedure are specific for File Gateways.

To prevent data loss in File Gateways that use ephemeral disks

1. Stop all the processes that are writing to Amazon S3.

2. Subscribe to receive notification from CloudWatch Events. For information, see [Getting notified about file operations](#).
3. Call the [NotifyWhenUploaded API](#) to get notified when data that is written, up until the ephemeral storage was lost, has been durably stored in Amazon S3.
4. Wait for the API to complete and you receive a notification id.

You receive a CloudWatch event with the same notification id.

5. Verify that the CachePercentDirty metric for your file share is 0. This confirms that all your data has been written to Amazon S3. For information about file share metrics, see [Understanding file share metrics](#).
6. You can now restart or stop the File Gateway without risk of losing any data.

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

- [Setting up your AWS Storage Gateway Hardware Appliance](#)
- [Physically installing your hardware appliance](#)
- [Accessing the hardware appliance console](#)
- [Configuring hardware appliance network parameters](#)
- [Activating your AWS Storage Gateway Hardware Appliance](#)

- [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

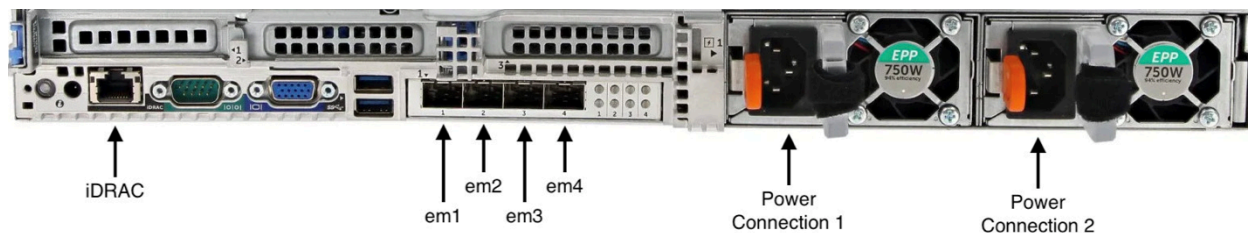
Before you perform the following procedure, make sure that you meet all of the requirements for the Storage Gateway Hardware Appliance as described in [Networking and firewall requirements for the Storage Gateway Hardware Appliance](#).

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.

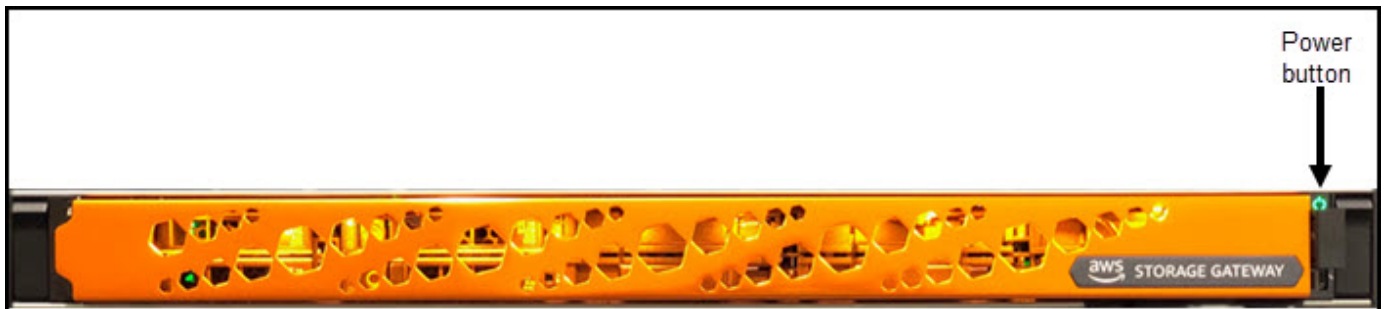
2. Plug in a power connection to each of the two power supplies. It's possible to plug in to only one power connection, but we recommend power connections to both power supplies for redundancy.

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 `Ctrl+]` (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 S3 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

- Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
- 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.
- 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 and activate an Amazon S3 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 S3 File Gateway](#)

- [Connect your Amazon S3 File Gateway to AWS](#)
- [Review settings and activate your Amazon S3 File Gateway](#)
- [Configure your Amazon S3 File Gateway](#)

Set up an Amazon S3 File Gateway

To set up a new S3 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 S3 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 that 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 S3 File Gateway to AWS

To connect a new S3 File Gateway to AWS

1. If you have not done so already, complete the procedure described in [Set up an Amazon S3 File Gateway](#). When finished, choose **Next** to open the **Connect to AWS** page in the AWS Storage Gateway console.
2. 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. For more information, see [Getting the gateway IP address](#)".

- **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.
3. In the **Endpoint options** section, for **Service endpoint**, choose the type of endpoint that 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.

Note

To specify a VPC endpoint that belongs to an AWS account other than the one you are currently using to create your gateway, you must provide its DNS name or IP address.

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 S3 File Gateway

To review settings and activate a new S3 File Gateway

1. If you have not done so already, complete the procedures described in the following topics:
 - [Set up an Amazon S3 File Gateway](#)
 - [Connect your Amazon S3 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 S3 File Gateway

To perform the first-time configuration on a new S3 File Gateway

1. If you have not done so already, complete the procedures described in the following topics:
 - [Set up an Amazon S3 File Gateway](#)
 - [Connect your Amazon S3 File Gateway to AWS](#)
 - [Review settings and activate your Amazon S3 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 about 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 notify you 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** – Don't receive CloudWatch notifications 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 create a file share for it to use. For instructions, see [Create a file share](#).

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 S3 File Gateway to AWS](#).

To configure your S3 File Gateway to transfer data through the VPC, you must create a separate VPC endpoint for Amazon S3, then specify this VPC endpoint when you create file shares for the gateway.

 **Note**

You must activate your gateway in the same region where you create the VPC endpoint for Storage Gateway, and the Amazon S3 storage that you configure for the file share must be in the same region where you create the VPC endpoint for Amazon S3.

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 S3 File Gateway](#).

For information about getting an activation key, see [Getting an activation key for your gateway](#).

Important

To configure your S3 File Gateway to transfer data through the VPC, you must create a separate VPC endpoint for Amazon S3, then specify this VPC endpoint when you create file shares for the gateway.

To do this, follow the same steps as shown above, but choose `com.amazonaws.region.s3` for **Service Name**, then select the route table that you want the S3 endpoint associated with instead of subnet/security group. For instructions, see [Creating a gateway endpoint](#).

Creating a file share

In this section, you can find instructions on how to create a file share that can be accessed using the Network File System (NFS) or the Server Message Block (SMB) protocol.

When you create an NFS share, anyone who has access to the NFS server can access the NFS file share by default. You can limit access to clients by IP address.

When you create an SMB file share, you can use one of three modes of authentication:

- A file share with Microsoft Active Directory (AD) access. Any authenticated Microsoft AD user gets access to this file share type.
- An SMB file share with limited access. Only certain domain users and groups that you specify are allowed access (through an allow list). Users and groups can also be denied access (through a deny list).
- An SMB file share with guest access. Any user who can provide the guest password has access to this file share.

Note

File shares that are exported through the gateway for NFS file shares support POSIX permissions. For SMB file shares, you can use access control lists (ACLs) to manage permissions on files and folders in your file share. For more information, see [Using Windows ACLs to limit SMB file share access](#).

A File Gateway can host one or more file shares of different types. You can have multiple NFS and SMB file shares on a File Gateway.

Important

To create a file share, a File Gateway requires you to activate AWS Security Token Service (AWS STS). If AWS STS isn't activated in the AWS Region where you create your File Gateway, activate it. For information about how to activate AWS STS, see [Activating and deactivating AWS Security Token Service in an AWS Region](#) in the *AWS Identity and Access Management User Guide*.

Topics

- [Avoiding unanticipated costs when uploading gateway data](#)
- [Encrypt objects stored by File Gateway in Amazon S3](#)
- [Create an NFS file share](#)
- [Create an SMB file share](#)

Avoiding unanticipated costs when uploading gateway data

When a file is written to the File Gateway by an NFS client, the File Gateway uploads the file's data to Amazon S3 followed by its metadata. Uploading the file data creates an S3 object, and uploading the metadata for the file updates the metadata for the S3 object. This process creates an additional version of the object. If S3 versioning is turned on, both versions are stored.

If you change the metadata of a file that's stored in your File Gateway, a new S3 object is created and replaces the existing S3 object. This behavior is different from editing a file in a file system, where editing a file does not result in a new file being created. Test all file operations that you plan to use with AWS Storage Gateway so that you understand how each file operation interacts with Amazon S3 storage.

Carefully consider the use of S3 versioning and Cross-Region replication (CRR) in Amazon S3 when you're uploading data from your File Gateway. Uploading files from your File Gateway to Amazon S3 when S3 versioning is turned on results in at least two versions of an S3 object.

Certain workflows involving large files and file-writing patterns such as file uploads that are performed in several steps can increase the number of stored S3 object versions. If the File Gateway cache needs to free up space due to high file-write rates, multiple S3 object versions might be created. These scenarios increase S3 storage if S3 Versioning is turned on and increase the transfer costs associated with CRR. Test all file operations that you plan to use with Storage Gateway so that you understand how each file operation interacts with Amazon S3 storage.

Using the Rsync utility with your File Gateway results in the creation of temporary files in the cache and the creation of temporary S3 objects in Amazon S3. This situation results in early deletion charges in the S3 Standard-Infrequent Access (S3 Standard-IA) storage class.

Encrypt objects stored by File Gateway in Amazon S3

S3 File Gateway supports the following methods of server-side encryption for the data that it stores in Amazon S3:

- **SSE-S3** — By default, all new objects uploaded to Amazon S3 buckets use server-side encryption with Amazon S3 managed keys. For more information, see [Using server-side encryption with Amazon S3 managed keys](#) in the *Amazon Simple Storage Service User Guide*.
- **SSE-KMS** — You can configure your file share to use server-side encryption with AWS Key Management Service (AWS KMS) managed keys. AWS KMS is a service that combines secure, highly available hardware and software to provide a key management system scaled for the cloud. For more information, see [What is AWS Key Management Service?](#) in the *AWS Key Management Service Developer Guide*.
- **DSSE-KMS** — Dual-layer server-side encryption with AWS KMS keys applies two layers of encryption to objects when they are uploaded to Amazon S3. This helps fulfill compliance standards for multilayer encryption. For more information, see [Using dual-layer server-side encryption with AWS KMS keys](#) in the *Amazon Simple Storage Service User Guide*.

Note

There are additional charges for using DSSE-KMS and AWS KMS keys. For more information, see [AWS KMS pricing](#).

You can specify an encryption method when you create a new file share by using the Storage Gateway console or the Storage Gateway API. For console procedures, see [Create an NFS file share with a custom configuration](#) or [Create an SMB file share with a custom configuration](#). For information about the corresponding API commands, see [CreateNFSFileShare](#) or [CreateSMBFileShare](#) in the *AWS Storage Gateway API Reference*.

You can also update encryption settings for an existing file share using the Storage Gateway console, or the Storage Gateway API. For the console procedure, see [Change the server-side encryption method for an existing file share](#). For information about the corresponding API commands, see [UpdateNFSFileShare](#) or [UpdateSMBFileShare](#) in the *AWS Storage Gateway API Reference*.

Note

After you update the encryption method, the gateway uses the new method for all new objects it creates in Amazon S3 and for any stored objects that it updates or modifies in the future. Existing Amazon S3 objects will only receive the new encryption method if they are updated or modified by the gateway.

Important

Make sure that your file share uses the same encryption type as the Amazon S3 bucket where it stores your data.

If you configure your File Gateway to use SSE-KMS or DSSE-KMS for encryption, you must manually add `kms:Encrypt`, `kms:Decrypt`, `kms:ReEncrypt*`, `kms:GenerateDataKey`, and `kms:DescribeKey` permissions to the IAM role associated with the file share. For more information, see [Using Identity-Based Policies \(IAM Policies\) for Storage Gateway](#).

Create an NFS file share

The Network File System (NFS) protocol is a stateful file sharing protocol for Unix-based systems. When an NFS-enabled client and NFS server communicate, the client requests a file or directory from the server using remote procedure calls (RPC). The server verifies that the file or directory is available and that the client has the required access permissions. The server then mounts the file or directory remotely on the client and shares access via a virtual connection. For client operations, NFS makes using the remote server file similar to accessing a local file.

Note

The NFS protocol supports a maximum of 16 groups per user. Users might have issues mounting NFS file shares if they belong to more than 16 groups. To avoid mounting issues, make sure that users are members of 16 or fewer groups when accessing NFS file shares.

The following topics explain various methods for creating an NFS file share for your File Gateway:

Contents

- [Create an NFS file share using the default configuration](#)
 - [Default configuration settings for NFS file shares](#)
- [Create an NFS file share with a custom configuration](#)

Create an NFS file share using the default configuration

This section explains how to create a new Network File System (NFS) file share using preconfigured default settings. Use this method for basic deployments, personal use, testing, or as a way to quickly deploy multiple file shares that you plan to edit and customize later. For a list of the default settings for file shares that you create using this procedure, see [Default configuration settings for NFS file shares](#). If you need more granular control or want to use advanced settings for your file share, see [Create an NFS file share using a custom configuration](#).

Note

If you need to connect your file share to Amazon S3 through a Virtual Private Cloud (VPC), you must follow the custom configuration procedure. You can't edit VPC settings for a file share after you create it.

Important

Using S3 Versioning, Cross-Region Replication, or the Rsync utility when uploading data from a File Gateway can have significant cost implications. For more information, see [Avoiding unanticipated costs when uploading data from File Gateway](#).

To create an NFS file share using the default configuration:

1. Open the AWS Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home/> and choose **File shares** from the left navigation pane.
2. Choose **Create file share**.
3. For **Gateway**, choose your Amazon S3 File Gateway from the list.
4. For **File share protocol**, choose **NFS**.
5. For **S3 bucket**, do one of the following:

- Choose an existing Amazon S3 bucket in your account from the dropdown list.
- Choose **A bucket in another account** from the dropdown list, then enter the name of the bucket in **Cross-account bucket name**.
- Choose **Create new S3 bucket**, then choose the AWS Region where the Amazon S3 endpoint for your new bucket is located, and enter a unique **S3 bucket name**. Choose **Create S3 bucket** when finished.

For information about creating a new bucket, see [How do I create an S3 bucket?](#) in the Amazon S3 User Guide.

 **Note**

S3 File Gateway does not support support Amazon S3 buckets with periods (.) in the bucket name.

Make sure your bucket name complies with the rules for bucket naming in Amazon S3. For more information, see [Rules for bucket naming](#) in the *Amazon Simple Storage Service User Guide*.

6. Review the settings under **Default configuration**, then choose **Create file share** to create your new NFS file share using the default configuration.

After your NFS file share is created, you can view its configuration settings in the AWS Storage Gateway console on the file share's **Details** tab. For information about mounting your file share, see [Mount your NFS file share on your client](#).

Default configuration settings for NFS file shares

The following settings apply to all new NFS file shares that you create using the default configuration. After you create a file share, you can select it from the **File shares** page in the AWS Storage Gateway console to view details about its configuration.

 **Important**

The default NFS file share configuration provides full file control and access permissions to the owner of the S3 bucket that's mapped to the file share, even if the bucket is owned by

a different AWS account. For more information about using your file share to access objects in a bucket that's owned by another account, see [Using a file share for cross-account access](#).

Setting	Default value	Notes
Amazon S3 location	The file share connects directly to the Amazon S3 bucket and has the same name as the bucket. Your gateway uses this bucket to store and retrieve files.	The name doesn't include a prefix.
AWS PrivateLink for S3	The file share doesn't connect to Amazon S3 through an interface endpoint in your virtual private cloud (VPC).	
File upload notification	Off	
Storage class for new objects	Amazon S3 Standard	This lets you store your frequently accessed object data redundantly in multiple Availability Zones that are geographically separated. For more information about the Amazon S3 Standard storage class, see Storage classes for frequently accessed objects in the <i>Amazon Simple Storage Service User Guide</i> .
Encryption	Server-side encryption with S3 managed keys (SSE-S3)	All Amazon S3 objects that your S3 File Gateway uploads, updates, or modifies are encrypted by default with

Setting	Default value	Notes
		server-side encryption using Amazon S3 managed keys.
Object metadata	Guess MIME type	<p>This allows Storage Gateway to guess the Multipurpose Internet Mail Extension (MIME) type for uploaded objects based on file extensions.</p> <p>This option requires that Access Control Lists (ACLs) are turned on for the Amazon S3 bucket that's associated with your file share. If ACLs are turned off, the file share can't access the Amazon S3 bucket, and remains in the Unavailable state indefinitely.</p>
Enable requester pays	Off	For more information, see Requester Pays buckets .
Audit logs	Off	Logging to an Amazon CloudWatch group is turned off by default.

Setting	Default value	Notes
Access to your S3 bucket	Create a new IAM role	The default option allows the File Gateway to create a new IAM role and access policy on your behalf. All NFS clients are allowed access. For information about supported NFS clients, see Supported NFS and SMB clients for File Gateway .
Mount options	<ul style="list-style-type: none"> Squash level – Root squash Export as – Read-write 	The default value of Squash level means that access for the remote superuser (root) is mapped to User Identifier (UID) (65534) and Group Identifier (GID) (65534).
File metadata defaults	<ul style="list-style-type: none"> Directory permissions – 0777 File permissions – 0666 User Identifier (UID) – 65534 Group Identifier (GID) – 65534 	

Create an NFS file share with a custom configuration

Use the following procedure to create a Network File System (NFS) file share with a custom configuration. To create an NFS file share using default configuration settings, see [Create an NFS file share using the default configuration](#).

⚠ Important

Using S3 Versioning, Cross-Region Replication, or the Rsync utility when uploading data from a File Gateway can have significant cost implications. For more information, see [Avoiding unanticipated costs when uploading data from File Gateway](#).

To create an NFS file share with customized settings

1. Open the AWS Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home/> and choose **File shares** from the left navigation pane.
2. Choose **Create file share**.
3. Choose **Customize configuration**. You can ignore the other fields on this page for now. You will be prompted to configure gateway, protocol, and storage settings in subsequent steps.
4. For **Gateway**, choose the Amazon S3 File Gateway for your new file share for from the dropdown list.
5. For **CloudWatch log group**, choose one of the following from the dropdown list:
 - To turn off logging for this file share, choose **Disable logging**.
 - To automatically create a new log group for this file share, choose **Created by Storage Gateway**.
 - To send health and resource notifications for this file share to an existing log group, choose the desired group from the list.

For more information about audit logs, see [Understanding S3 File Gateway audit logs](#).

6. (Optional) Under **Tags - Optional**, choose **Add new tag**, then enter a **Key** and **Value** for your file share.

A tag is a case-sensitive key-value pair that helps you categorize your Storage Gateway resources. Adding tags can make filtering and searching for your file share easier. You can repeat this step to add up to 50 tags.

Choose **Next** when finished.

7. For **S3 bucket**, do one of the following to specify where your file share will store and retrieve files:

- To connect the file share directly to an existing S3 bucket in your Amazon Web Services account, choose the bucket name from the dropdown list.
- To connect the file share to an existing S3 bucket that is owned by an Amazon Web Services account other than the one that you use to create the file share, choose **A bucket in another account** from the dropdown list, then enter the **Cross-account bucket name**.
- To connect the file share to a new S3 bucket, choose **Create a new S3 bucket**, then choose the **Region** where the Amazon S3 endpoint for your new bucket is located, and enter a unique **S3 bucket name**. Choose **Create S3 bucket** when finished. For more information about creating new buckets, see [How do I create an S3 bucket?](#) in the Amazon S3 User Guide.
- To connect the file share to an S3 bucket using an access point name, choose **Amazon S3 access point name** from the dropdown list, then enter the **Access point name**. If you need to create a new access point, you can choose **Create an S3 access point**. For further instructions, see [Creating an access point](#) in the Amazon S3 User Guide. For more information about access points, see [Managing data access with Amazon S3 access points](#) and [Delegating access control to access points](#) in the Amazon S3 User Guide.
- To connect the file share to an S3 bucket using an access point alias, choose **Amazon S3 access point alias** from the dropdown list, then enter the **Access point alias**. If you need to create a new access point, you can choose **Create an S3 access point**. For further instructions, see [Creating an access point](#) in the Amazon S3 User Guide. For more information about access point aliases, see [Using a bucket-style alias for your access point](#) in the Amazon S3 User Guide.

Note

Each file share can only connect to one S3 bucket, but multiple file shares can connect to the same bucket. If you connect more than one file share to the same bucket, you must configure each file share to use a unique, non-overlapping **S3 bucket prefix** to prevent read/write conflicts.

S3 File Gateway does not support support Amazon S3 buckets with periods (.) in the bucket name.

Make sure your bucket name complies with the rules for bucket naming in Amazon S3. For more information, see [Rules for bucket naming](#) in the *Amazon Simple Storage Service User Guide*.

8. (Optional) For **S3 bucket prefix**, enter a prefix for your file share to apply to the objects it creates in Amazon S3. Prefixes are a way to organize your data in S3, similar to directories in traditional file structures. For more information, see [Organizing objects using prefixes](#) in the Amazon S3 User Guide.

 **Note**

- If you connect more than one file share to the same bucket, you must configure each file share to use a unique, non-overlapping prefix to prevent read/write conflicts.
- The prefix must end with a forward slash (/).
- After the file share is created, the prefix can't be modified or deleted.

9. For **Region**, choose the AWS Region where the S3 endpoint for your bucket is located from the dropdown list. This field appears only when you specify an access point or a bucket in another account for **S3 bucket**.
10. For **Storage class for new objects**, choose a storage class from the dropdown list. For more information about storage classes, see [Using storage classes with a File Gateway](#).
11. For **IAM Role**, do one of the following to configure an IAM role for your file share:
 - To automatically create a new IAM role with the necessary permissions for your file share to work properly, choose **Created by Storage Gateway** from the dropdown list.
 - To use an existing IAM role, choose the role name from the dropdown list.
 - To create a new IAM role, choose **Create a role**. For further instructions, see [Creating a role to delegate permissions to an AWS service](#) in the AWS Identity and Access Management User Guide.

For more information about how IAM roles control access between your file share and S3 bucket, see [Granting access to an Amazon S3 bucket](#).

12. For **Private link**, do the following only if you need to configure your file share to communicate with AWS using a private endpoint in a Virtual Private Cloud (VPC). Otherwise, skip this step. For more information, see [What is AWS PrivateLink?](#) in the AWS PrivateLink Guide.
 - a. Select **Use VPC endpoint**.
 - b. For **Identify VPC endpoint by**, do one of the following:

- Select **VPC endpoint ID**, then choose the endpoint that you want to use from the **VPC endpoint** dropdown list.
- Select **DNS name**, then enter the **DNS name** for the endpoint that you want to use.

13. For **Encryption**, choose the type of server-side encryption that the file share will use for the data that it stores in Amazon S3:

- To use server-side encryption managed with Amazon S3 (SSE-S3), choose **S3-Managed Keys (SSE-S3)**.

For more information, see [Using server-side encryption with Amazon S3 managed keys](#) in the *Amazon Simple Storage Service User Guide*.

- To use server-side encryption managed with AWS Key Management Service (SSE-KMS), choose **KMS-Managed Keys (SSE-KMS)**. For **Primary KMS key**, choose an existing AWS KMS key, or choose **Create a new KMS key** to create a new KMS key in the AWS Key Management Service (AWS KMS) console.

For more information about AWS KMS, see [What is AWS Key Management Service?](#) in the *AWS Key Management Service Developer Guide*.

- To use dual-layer server-side encryption managed with AWS Key Management Service (DSSE-KMS), choose **Dual-layer server-side encryption with AWS Key Management Service keys (DSSE-KMS)**. For **Primary KMS key**, choose an existing AWS KMS key, or choose **Create a new KMS key** to create a new KMS key in the AWS Key Management Service (AWS KMS) console.

For more information about DSSE-KMS, see [Using dual-layer server-side encryption with AWS KMS keys](#) in the *Amazon Simple Storage Service User Guide*.

 **Note**

There are additional charges for using DSSE-KMS and AWS KMS keys. For more information, see [AWS KMS pricing](#).

To specify an AWS KMS key with an alias that is not listed or to use an AWS KMS key from a different AWS account, you must use the AWS Command Line Interface. Asymmetric KMS keys are not supported. For more information, see [CreateNFSFileShare](#) in the *AWS Storage Gateway API Reference*.

 **Important**

Make sure that your file share uses the same encryption type as the Amazon S3 bucket where it stores your data.

14. For **Guess MIME types**, select **Guess media MIME type** to allow Storage Gateway to guess the media type for uploaded objects based on their file extensions.
15. For **File share name**, enter a name for your file share.

 **Note**

A valid NFS file share name can only contain the following characters: a-z, A-Z, 0-9, -, ., and _.

16. For **Upload events**, select **Log an event when a file is successfully uploaded by the gateway** if you want your gateway to record CloudWatch log events when it successfully uploads files to Amazon S3. Notification delay controls the minimum delay between the most recent client write operation and generation of the ObjectUploaded log notification. Because clients can make many small writes to files in a short time, we recommend setting this parameter for as long as possible to avoid generating multiple notifications for the same file in rapid succession. For more information, see [Getting file upload notification](#).

 **Note**

This setting has no effect on the timing of the object uploading to S3, only on the timing of the notification.

This setting is not meant to specify an exact time at which the notification will be sent. In some cases, the gateway might require more than the specified delay time to generate and send notifications.

Choose **Next** when finished.

- 17.
18. For **File share protocol**, choose **NFS**.

19. For **Client access**, do one of the following to specify which NFS clients can access your file share:

- To accept all incoming client connections, select **All NFS clients**.
- To accept incoming client connections only from specific IP addresses, select **Specific NFS clients**, then choose **Add a client**. For **Allowed clients**, specify a valid IP address or CIDR block from which to accept connections. If you need to specify additional IP addresses, choose **Add another client**.

 **Note**

We recommend configuring limiting access to your file share using the **Specific NFS clients** option. If you don't, any client on your network can mount to the file share.

20. For **Access type**, select one of the following:

- To allow clients to read and write files on the file share, select **Read/Write**.
- To allow clients to read files but not write to the file share, select **Read-only**.

 **Note**

For file shares that are mounted on a Microsoft Windows client, if you choose **Read-only**, you might see a message about an unexpected error keeping you from creating the folder. You can ignore this message.

21. For **Access level**, choose one of the following:

- **Root squash (default)**: Access for the remote superuser (root) is mapped to UID (65534) and GID (65534).
- **All squash**: All user access is mapped to User ID (UID) (65534) and Group ID (GID) (65534).
- **No root squash**: The remote superuser (root) receives access as root.

22. (Optional) For **Automated cache refresh from S3**, choose **Set cache refresh interval**, then set the time in **Minutes** or **Days** to refresh the file share's cache using Time To Live (TTL). TTL is the length of time since the last refresh. After the TTL interval has elapsed, accessing a directory causes the File Gateway to refresh that directory's contents from the Amazon S3 bucket.

 **Note**

Setting this value shorter than 30 minutes can negatively impact gateway performance in situations where large numbers of Amazon S3 objects are frequently created or deleted.

23. For **File metadata defaults**, select **Change default metadata for S3 objects that were not created or modified by your gateway** if you want your gateway to apply file metadata (including Unix permissions) to preexisting objects that it discovers in your S3 bucket. Specify the **Directory permissions**, **File permissions**, **User ID**, and **Group ID** that you want to apply in the corresponding fields.
24. For **File ownership and permissions**, select **Give the S3 bucket owner full ownership of files created by the gateway, including read, write, edit, and delete permissions** if you want the AWS account that owns the S3 bucket to have full control of all objects written to the bucket by your file share.

Choose **Next** when finished.

25. Review the file share configuration. Choose **Edit** to modify the settings for any section that you want to change. When finished, choose **Create**.

After your NFS file share is created, you can view its configuration settings in the AWS Storage Gateway console on the file share's **Details** tab. For instructions to mount your file share, see [Mount your NFS file share on your client](#).

Create an SMB file share

The Server Message Block (SMB) protocol is deeply integrated into the Microsoft Windows product suite, and remains the default file sharing protocol for Windows operating systems. The process of client-server communication is similar to NFS at a high level, but there are differences in some details and operational mechanisms. For example, in SMB, file systems are not mounted on the local SMB client. Instead, a network share hosted on the SMB server is accessed via a network path.

The topics in this section explain various methods for creating an SMB file share for your File Gateway.

Contents

- [Create an SMB file share using the default configuration](#)
 - [Default configuration settings for SMB file shares](#)
- [Create an SMB file share with a custom configuration](#)

Create an SMB file share using the default configuration

This section explains how to create a new Server Message Block (SMB) file share using preconfigured default settings. Use this method for basic deployments, personal use, testing, or as a way to quickly deploy multiple file shares that you plan to edit and customize later. For a list of the default settings for file shares that you create using this procedure, see [Default configuration settings for SMB file shares](#). If you need more granular control or want to use advanced settings for your file share, see [Create an SMB file share with a custom configuration](#).

Note

If you need to connect your file share to Amazon S3 through a Virtual Private Cloud (VPC), you must follow the custom configuration procedure. You can't edit VPC settings for a file share after you create it.

Important

Using S3 Versioning, Cross-Region Replication, or the Rsync utility when uploading data from a File Gateway can have significant cost implications. For more information, see [Avoiding unanticipated costs when uploading data from File Gateway](#).

Prerequisites

Before you create your file share, do the following:

- Configure SMB security settings for your File Gateway. For instructions, see [Setting a security level for your gateway](#).
- Configure either Microsoft Active Directory or guest access for authentication. For instructions, see [Using Active Directory to authenticate users](#) or [Providing guest access to your file share](#).
- Make sure that the required ports are open in your security group. For more information, see [Port Requirements](#).

To create an SMB file share using the default configuration:

1. Open the AWS Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home/> and choose **File shares** from the left navigation pane.
2. Choose **Create file share**.
3. For **Gateway**, choose the Amazon S3 File Gateway from the dropdown list.
4. For **File share protocol**, choose **SMB**.
5. For **S3 bucket**, do one of the following:
 - Choose an existing Amazon S3 bucket in your account from the dropdown list.
 - Choose **A bucket in another account** from the dropdown list, then enter the name of the bucket in **Cross-account bucket name**.
 - Choose **Create new S3 bucket**, then choose the AWS Region where the Amazon S3 endpoint for your new bucket is located, and enter a unique **S3 bucket name**. Choose **Create S3 bucket** when finished.

For information about creating a new bucket, see [How do I create an S3 bucket?](#) in the Amazon S3 User Guide.

Note

S3 File Gateway does not support support Amazon S3 buckets with periods (.) in the bucket name.

Make sure your bucket name complies with the rules for bucket naming in Amazon S3. For more information, see [Rules for bucket naming](#) in the *Amazon Simple Storage Service User Guide*.

6. **User authentication**, choose the authentication method you want to use from the dropdown list:
 - To use your corporate Microsoft Active Directory or AWS Managed Microsoft AD to authenticate user access to your SMB file share, choose **Active Directory**. Your gateway must be joined to a domain to use this method. For more information, see [Using Active Directory to authenticate users](#).

 **Note**

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](#).

If **Join status** indicates that your gateway is already joined to an Active Directory domain, proceed to the next step. Otherwise, do the following:

1. Choose **Configure**.
2. For **Domain**, enter the name of the Active Directory domain you want your gateway to join.
3. Enter the **Username** and **Password** that the gateway will use to join the domain.
4. (Optional) For **Organization unit (OU)**, enter the designated OU that your Active Directory uses for new computer objects.
5. (Optional) For **Domain controller(s) (DC)**, enter the name of the DC through which your gateway will connect to Active Directory. You can leave this field blank to allow DNS to automatically select a DC.
6. Choose **Join Active Directory**.

 **Note**

Joining a domain creates an Active Directory account in the default container (which isn't an organizational unit) using the 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 domain join process, you 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.

- To grant password-protected access to anyone who provides the guest password that you configure, choose **Guest access**. Your File Gateway doesn't need to be part of a Microsoft Active Directory domain to use this method. Choose **Configure** to specify your **Guest password**, then choose **Save**.
7. Review the settings under **Default configuration**, then choose **Create file share** to create your new SMB file share using the default configuration.

After your SMB file share is created, you can view its configuration settings in the AWS Storage Gateway console on the file share's **Details** tab. For information about mounting your file share, see [Mount your SMB file share on your client](#).

Default configuration settings for SMB file shares

The following settings apply to all new SMB file shares that you create using the default configuration. After you create a file share, you can select it from the **File shares** page in the AWS Storage Gateway console to view details about its configuration.

Important

The default SMB file share configuration provides full file control and access permissions to the owner of the S3 bucket that's mapped to the file share, even if the bucket is owned by a different Amazon Web Services account. For more information about using your file share to access objects in a bucket that's owned by another account, see [Using a file share for cross-account access](#).

Setting	Default value	Notes
Amazon S3 location	The file share connects directly to the Amazon S3 bucket and has the same name as the bucket. Your gateway uses this bucket to store and retrieve files.	The name doesn't include a prefix.
AWS PrivateLink for S3	The file share doesn't connect to Amazon S3 through an	

Setting	Default value	Notes
	interface endpoint in your virtual private cloud (VPC).	
File upload notification	Off	
Storage class for new objects	Amazon S3 Standard	This lets you store your frequently accessed object data redundantly in multiple Availability Zones that are geographically separated. For more information about the Amazon S3 Standard storage class, see Storage classes for frequently accessed objects in the <i>Amazon Simple Storage Service User Guide</i> .
Encryption	Server-side encryption with S3 managed keys (SSE-S3)	All Amazon S3 objects that your S3 File Gateway uploads, updates, or modifies are encrypted by default with server-side encryption using Amazon S3 managed keys.

Setting	Default value	Notes
Object metadata	Guess MIME type	<p>This allows Storage Gateway to guess the Multipurpose Internet Mail Extension (MIME) type for uploaded objects based on file extensions.</p> <p>This option requires that Access Control Lists (ACLs) are turned on for the Amazon S3 bucket that's associated with your file share. If ACLs are turned off, the file share can't access the Amazon S3 bucket, and remains in the Unavailable state indefinitely.</p>
Access based enumeration	Not activated	<p>The files and folders on the file share are visible to all users during directory enumeration. Access-based enumeration is a system that filters the enumeration of files and folders on an SMB file share based on the share's access control lists (ACLs).</p>
Enable requester pays	Off	<p>For more information, see Requester Pays buckets.</p>

Setting	Default value	Notes
Opportunistic locking	On	This allows the file share to use opportunistic locking to optimize the file buffering strategy. In most cases, activating opportunistic locking improves performance, particularly with regard to Windows context menus.
Audit logs	Off	Logging to an Amazon CloudWatch group is turned off by default.
Force case sensitivity	Off	This allows the client to control the case sensitivity.
Access to your S3 bucket	Create a new IAM role	The default option allows the File Gateway to create a new IAM role and access policy on your behalf.

Create an SMB file share with a custom configuration

Use the following procedure to create a Server Message Block (SMB) file share with a custom configuration. To create an SMB file share using default configuration settings, see [Create an SMB file share using the default configuration](#).

Important

Using S3 Versioning, Cross-Region Replication, or the Rsync utility when uploading data from a File Gateway can have significant cost implications. For more information, see [Avoiding unanticipated costs when uploading data from File Gateway](#).

Prerequisites

Before you create your file share, do the following:

- Configure SMB security settings for your File Gateway. For instructions, see [Setting a security level for your gateway](#).
- Configure either Microsoft Active Directory or guest access for authentication. For instructions, see [Using Active Directory to authenticate users](#) or [Providing guest access to your file share](#).
- Make sure that the required ports are open in your security group. For more information, see [Port Requirements](#).

To create an SMB file share with customized settings

1. Open the AWS Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home/> and choose **File shares** from the left navigation pane.
2. Choose **Create file share**.
3. Choose **Customize configuration**. You can ignore the other fields on this page for now. You will be prompted to configure gateway, protocol, and storage settings in subsequent steps.
4. For **Gateway**, choose the Amazon S3 File Gateway from the dropdown list.
5. For **CloudWatch log group**, choose one of the following from the dropdown list:
 - To turn off logging for this file share, choose **Disable logging**.
 - To automatically create a new log group for this file share, choose **Created by Storage Gateway**.
 - To send health and resource notifications for this file share to an existing log group, choose the desired group from the list.

For more information about audit logs, see [Understanding S3 File Gateway audit logs](#).

6. (Optional) Under **Tags - Optional**, choose **Add new tag**, then enter a **Key** and **Value** for your file share. A tag is a case-sensitive key-value pair that helps you to categorize your Storage Gateway resources. Adding tags can make filtering and searching for your file share easier. You can repeat this step to add up to 50 tags.

Choose **Next** when finished.

7. For **S3 bucket**, do one of the following to specify where to store and retrieve files:
 - To connect the file share directly to an existing S3 bucket in your Amazon Web Services account, choose the bucket name from the dropdown list.
 - To connect the file share to an existing S3 bucket that's owned by an Amazon Web Services account other than the one that you're using to create the file share, choose **A bucket in another account** from the dropdown list, then enter the **Cross-account bucket name**.
 - To connect the file share to a new S3 bucket, choose **Create a new S3 bucket**, then choose the **Region** where the Amazon S3 endpoint for your new bucket is located, and enter a unique **S3 bucket name**. Choose **Create S3 bucket** when finished. For more information about creating new buckets, see [How do I create an S3 bucket?](#) in the Amazon S3 User Guide.
 - To connect the file share to an S3 bucket using an access point name, choose **Amazon S3 access point name** from the dropdown list, then enter the **Access point name**. If you need to create a new access point, you can choose **Create an S3 access point**. For further instructions, see [Creating an access point](#) in the Amazon S3 User Guide. For more information about access points, see [Managing data access with Amazon S3 access points](#) and [Delegating access control to access points](#) in the Amazon S3 User Guide.
 - To connect the file share to an S3 bucket using an access point alias, choose **Amazon S3 access point alias** from the dropdown list, then enter the **Access point alias**. If you need to create a new access point, you can choose **Create an S3 access point**. For further instructions, see [Creating an access point](#) in the Amazon S3 User Guide. For more information about access point aliases, see [Using a bucket-style alias for your access point](#) in the Amazon S3 User Guide.

 **Note**

Each file share can only connect to one S3 bucket, but multiple file shares can connect to the same bucket. If you connect more than one file share to the same bucket, you

must configure each file share to use a unique, non-overlapping **S3 bucket prefix** to prevent read/write conflicts.

S3 File Gateway does not support support Amazon S3 buckets with periods (.) in the bucket name.

Make sure your bucket name complies with the rules for bucket naming in Amazon S3. For more information, see [Rules for bucket naming](#) in the *Amazon Simple Storage Service User Guide*.

8. (Optional) For **S3 bucket prefix**, enter a prefix for your file share to apply to the objects it creates in Amazon S3. Prefixes are a way to organize your data in S3, similar to directories in traditional file structures. For more information, see [Organizing objects using prefixes](#) in the Amazon S3 User Guide.

 **Note**

- If you connect more than one file share to the same bucket, you must configure each file share to use a unique, non-overlapping prefix to prevent read/write conflicts.
- The prefix must end with a forward slash (/).
- After the file share is created, the prefix can't be modified or deleted.

9. For **Region**, choose the AWS Region where the S3 endpoint for your bucket is located from the dropdown list. This field appears only when you specify an access point or a bucket in another account for **S3 bucket**.
10. For **Storage class for new objects**, choose a storage class from the dropdown list. For more information about storage classes, see [Using storage classes with a File Gateway](#).
11. For **IAM Role**, do one of the following to configure an IAM role for your file share:
 - To automatically create a new IAM role with the necessary permissions for your file share to work properly, choose **Created by Storage Gateway** from the dropdown list.
 - To use an existing IAM role, choose the role name from the dropdown list.
 - To create a new IAM role, choose **Create a role**. For further instructions, see [Creating a role to delegate permissions to an AWS service](#) in the AWS Identity and Access Management User Guide.

For more information about how IAM roles control access between your file share and S3 bucket, see [Granting access to an Amazon S3 bucket](#).

12. For **Private link**, do the following only if you need to configure your file share to communicate with AWS using a private endpoint in a Virtual Private Cloud (VPC). Otherwise, skip this step. For more information, see [What is AWS PrivateLink?](#) in the AWS PrivateLink Guide.

- a. Select **Use VPC endpoint**.
- b. For **Identify VPC endpoint by**, do one of the following:
 - Select **VPC endpoint ID**, then choose the endpoint that you want to use from the **VPC endpoint** dropdown list.
 - Select **DNS name**, then enter the **DNS name** for the endpoint that you want to use.

13. For **Encryption**, choose the type of encryption keys to use to encrypt objects that your File Gateway stores in Amazon S3:

- To use server-side encryption managed with Amazon S3 (SSE-S3), choose **S3-Managed Keys (SSE-S3)**.

For more information, see [Using server-side encryption with Amazon S3 managed keys](#) in the *Amazon Simple Storage Service User Guide*.

- To use server-side encryption managed with AWS Key Management Service (SSE-KMS), choose **KMS-Managed Keys (SSE-KMS)**. For **Primary KMS key**, choose an existing AWS KMS key, or choose **Create a new KMS key** to create a new KMS key in the AWS Key Management Service (AWS KMS) console.

For more information about AWS KMS, see [What is AWS Key Management Service?](#) in the *AWS Key Management Service Developer Guide*.

- To use dual-layer server-side encryption managed with AWS Key Management Service (DSSE-KMS), choose **Dual-layer server-side encryption with AWS Key Management Service keys (DSSE-KMS)**. For **Primary KMS key**, choose an existing AWS KMS key, or choose **Create a new KMS key** to create a new KMS key in the AWS Key Management Service (AWS KMS) console.

For more information about DSSE-KMS, see [Using dual-layer server-side encryption with AWS KMS keys](#) in the *Amazon Simple Storage Service User Guide*.

Note

There are additional charges for using DSSE-KMS and AWS KMS keys. For more information, see [AWS KMS pricing](#).

To specify an AWS KMS key with an alias that is not listed or to use an AWS KMS key from a different AWS account, you must use the AWS Command Line Interface. Asymmetric KMS keys are not supported. For more information, see [CreateSMBFileShare](#) in the *AWS Storage Gateway API Reference*.

Important

Make sure that your file share uses the same encryption type as the Amazon S3 bucket where it stores your data.

14. For **Guess MIME types**, select **Guess media MIME type** to allow Storage Gateway to guess the Multipurpose Internet Mail Extension (MIME) type for uploaded objects based on their file extensions.
15. For **File share name**, enter a name for your file share.

Note

A valid SMB file share name cannot contain the following characters:
[,],#,;,<,>,:,\",\\,/,,|,?,*,+, or ASCII control characters 1-31.

16. For **Upload events**, select **Log an event when a file is successfully uploaded by the gateway** if you want your gateway to record CloudWatch log events when it successfully uploads files to Amazon S3. Notification delay controls the delay between the most recent client write operation and generation of the ObjectUploaded log notification. Because clients can make many small writes to files in a short time, we recommend setting this parameter for as long as possible to avoid generating multiple notifications for the same file in rapid succession. For more information, see [Getting file upload notification](#).

Note

This setting has no effect on the timing of the object uploading to S3, only on the timing of the notification.

This setting is not meant to specify an exact time at which the notification will be sent. In some cases, the gateway might require more than the specified delay time to generate and send notifications.

Choose **Next** when finished.

17. For **File share protocol**, choose **SMB**.

18. For **User authentication**, choose the authentication method that you want to use from the dropdown list:

- To use your corporate Microsoft Active Directory or AWS Managed Microsoft AD to authenticate user access to your SMB file share, choose **Active Directory**. Your gateway must be joined to a domain to use this method. For more information, see [Using Active Directory to authenticate users](#).

 **Note**

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](#).

If **Join status** indicates that your gateway is already joined to an Active Directory domain, proceed to the next step. Otherwise, do the following:

1. Choose **Configure**.
2. For **Domain**, enter the name of the Active Directory domain that you want your gateway to join.
3. Enter the **Username** and **Password** that the gateway will use to join the domain.
4. (Optional) For **Organization unit (OU)**, enter the designated OU that your Active Directory uses for new computer objects.

5. (Optional) For **Domain controller(s) (DC)**, enter the name of the DC through which your gateway will connect to Active Directory. You can leave this field blank to allow DNS to automatically select a DC.
6. Choose **Join Active Directory**.

 **Note**

Joining a domain creates an Active Directory account in the default container (which isn't an organizational unit), 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 domain join process, you 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.

- To grant password-protected access to anyone who provides the guest password that you configure, choose **Guest access**. Your File Gateway doesn't need to be part of a Microsoft Active Directory domain to use this method. Choose **Configure** to specify your **Guest password**, then choose **Save**.
19. For **User access**, do one of the following to specify which SMB clients can access your file share:
- To grant access to all users that successfully authenticate through Active Directory, select **All AD-authenticated users**.
 - To allow or deny access to specific users or groups, choose **Specific AD-authenticated users or groups**, then do the following:
 - For **Allowed users and groups**, choose **Add allowed user** or **Add allowed group** and enter an Active Directory user or group that you want to allow file share access. Repeat this process to allow as many users and groups as necessary
 - For **Denied users and groups**, choose **Add denied user** or **Add denied group** and enter an Active Directory user or group that you want to deny file share access. Repeat this process to deny as many users and groups as necessary.

Note

The **User and group file share access** section appears only if User authentication is set to **Active Directory**.

When specifying users or groups, do not include the domain. The domain name is implied by the membership of the gateway in the specific Active Directory to which it is joined.

20. (Optional) For **Admin users**, enter a comma-separated list of Active Directory users and groups. Admin users receive privileges to update access control lists (ACLs) on all files and folders in the file share. Groups must be prefixed with the @ character, for example, @group1.
21. For **Access type**, select one of the following:
 - To allow clients to read and write files on the file share, select **Read/Write**.
 - To allow clients to read files but not write to the file share, select **Read-only**.

Note

For file shares that are mounted on a Microsoft Windows client, if you choose **Read-only**, you might see a message about an unexpected error keeping you from creating the folder. You can ignore this message.

22. For **File and directory access control**, select one of the following:
 - To set fine-grained permissions on files and folders in your SMB file share, select **Windows Access Control List**. For more information, see [Using Microsoft Windows ACLs to Control Access to an SMB File Share](#).
 - To use POSIX permissions to control access to files and directories that are stored through your SMB file share, choose **POSIX permissions**.
23. For **Access based enumeration**, do one of the following:
 - To make the files and folders on the share visible only to users who have read access, select **Hide files and directories where user doesn't have permission**.
 - To make the files and folders on the share visible to all users during directory enumeration, don't select the check box.

 **Note**

Access-based enumeration is a system that filters the enumeration of files and folders on an SMB file share based on the share's access control lists (ACLs).

24. For File access options, select one of the following:

- To optimize the file share's file buffering strategy using opportunistic locking, select **Opportunistic lock**. In most cases, activating opportunistic locking improves performance, particularly with regard to Windows context menus.
- To allow the gateway - rather than the SMB client - to control file name case sensitivity, select **Force case sensitivity**.
- To deactivate both settings, select **Neither**.

 **Note**

To avoid file access conflicts, these settings are mutually exclusive and cannot be activated at the same time.

25. (Optional) For **Automated cache refresh from S3**, choose **Set cache refresh interval**, then set the time in **Minutes** or **Days** to refresh the file share's cache using Time To Live (TTL). TTL is the length of time since the last refresh. After the TTL interval has elapsed, accessing a directory causes the File Gateway to refresh that directory's contents from the Amazon S3 bucket.

 **Note**

Setting this value shorter than 30 minutes can negatively impact gateway performance in situations where large numbers of Amazon S3 objects are frequently created or deleted.

26. For **File ownership and permissions**, select **Give the S3 bucket owner full ownership of files created by the gateway, including read, write, edit, and delete permissions** if you want the AWS account that owns the S3 bucket to have full control of all objects written to the bucket by your file share.

Choose **Next** when finished.

27. Review the file share configuration. Choose **Edit** to modify the settings for any section that you want to change. When finished, choose **Create**.

After your SMB file share is created, you can view its configuration settings in the AWS Storage Gateway console on the file share's **Details** tab. For instructions to mount your file share, see [Mount your SMB file share on your client](#).

Mounting and using your file share

The topics in this section provide instructions about how to mount your file share on your client, use your file share, test your File Gateway, and clean up resources that are no longer needed, such as the gateways, Amazon EC2 instances, or on-premises VMs that you might create for testing purposes. For more information about supported Network File System (NFS) and Service Message Block (SMB) clients, see [Supported NFS and SMB clients for File Gateway](#).

Note

The AWS Management Console also provides example commands that you can use to mount your file share.

Topics

- [Mount your NFS file share on your client](#) - Learn how to mount your NFS file share on a drive on your client and map it to your Amazon S3 bucket.
- [Mount your SMB file share on your client](#) - Learn how to mount your SMB file share and map to a drive accessible to your client.
- [Using file shares on buckets with pre-existing objects](#) - Learn how to export a file share on an Amazon S3 bucket with objects created outside of the File Gateway using either NFS or SMB.
- [Test your S3 File Gateway](#) - Learn how to test your gateway by copying files and folders to your mapped drive and verifying that they appear in your Amazon S3 bucket automatically.

Mount your NFS file share on your client

Use the following procedure to mount your NFS file share on a drive on your client and map it to your Amazon S3 bucket.

To mount a file share and map it to an Amazon S3 bucket

1. If you are using a Microsoft Windows client, we recommend that you [create an SMB file share](#) and access it using an SMB client that is already installed on Windows client. If you use NFS, turn on Services for NFS in Windows.
2. Mount your NFS file share:

- For Linux clients, type the following command at the command prompt.

```
sudo mount -t nfs -o nolock,hard [GatewayVMIPAddress]:/[FileShareName]  
[ClientMountPath]
```

- For Windows clients, type the following command at the command prompt (**cmd.exe**).

```
mount -o nolock -o mtype=hard [GatewayVMIPAddress]:/[FileShareName]  
[WindowsDriveLetter]
```

For example, suppose that on a Windows client your VM's IP address is 123.123.1.2 and your file share name is test-fileshare. Suppose also that you want to map to drive T. In this case, your command looks like the following.

```
mount -o nolock -o mtype=hard 123.123.1.2:/test-fileshare T:
```

Note

When mounting file shares, be aware of the following:

- By default, Windows uses a soft mount for NFS shares. Soft mounts time out more easily when connection issues occur. We recommend using a hard mount for critical workloads, because hard mounts are safer and better preserve your data. To use a hard mount, make sure your command uses the `-o mtype=hard` switch.
- S3 File Gateway does not support NFS file locking. Always use the `-o nolock` option to turn off file locking when mounting NFS file shares.
- You might have a case where a folder and an object exist in an Amazon S3 bucket and have the same name. In this case, if the object name doesn't contain a trailing slash, only the folder is visible in a File Gateway. For example, if a bucket contains an object named test or test/ and a folder named test/test1, only test/ and test/test1 are visible in a File Gateway.
- You might need to remount your file share after a reboot of your client.
- If you are using Windows clients, check your mount options after mounting by running the mount command with no options. The response should confirm the file share was mounted using the latest options you provided. It also should confirm that you are not using cached old entries, which take at least 60 seconds to clear.

Next Step

[Test your S3 File Gateway](#)

Mount your SMB file share on your client

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

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

- **Command Prompt (cmdkey and net use)** – Use the command prompt to mount your file share. Store your credentials with cmdkey, then mount the drive with net use and include the /persistent:yes and /savecred switches if you want the connection to persist across system reboots. The specific commands you use will be different depending on whether you want to mount the drive for Microsoft Active Directory (AD) access or guest user access. Examples are provided below.
- **File Explorer (Map Network Drive)** – Use Windows File Explorer to mount your file share. Configure settings to specify whether you want the connection to persist across system reboots and prompt for network credentials.
- **PowerShell script** – Create a custom PowerShell script to mount your file share. Depending on the parameters you specify in the script, the connection can be persistent across system reboots, and the share can be either visible or invisible to the operating system while mounted.

Note

If you are a Microsoft AD 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.

If you are a guest user, make sure that you have the guest user password before attempting to mount the file share.

To mount your SMB file share for authorized Microsoft AD users using the command prompt:

1. Make sure the Microsoft AD user has the necessary permissions to the SMB file share before mounting the file share to the user's system.

2. Enter the following at the command prompt to mount the file share:

```
net use WindowsDriveLetter: \\GatewayIPAddress\FileShareName /  
persistent:yes
```

To mount your SMB file share with a specific sign-in credentials combination using the command prompt:

1. Make sure that the user has access to the SMB file share before mounting the file share to the system.
2. Enter the following at the command prompt to save the user credentials in Windows Credential Manager:

```
cmdkey /add:GatewayIPAddress /user:DomainName\UserName /pass:Password
```

3. Enter the following at the command prompt to mount the file share:

```
net use WindowsDriveLetter: \\GatewayIPAddress\FileShareName /  
persistent:yes /savecred
```

To mount your SMB file share for guest users using the command prompt:

1. Make sure that you have the guest user password before mounting the file share.
2. Type the following at the command prompt to save the guest credentials in Windows Credential Manager:

```
cmdkey /add:GatewayIPAddress /user:DomainName\smbguest /pass:Password
```

3. Type the following at the command prompt.

```
net use WindowsDriveLetter: \\$GatewayIPAddress\$Path /user:$Gateway  
ID\smbguest /persistent:yes /savecred
```

 **Note**

When mounting file shares, be aware of the following:

- You might have a case where a folder and an object exist in an Amazon S3 bucket and have the same name. In this case, if the object name doesn't contain a trailing slash, only

the folder is visible in a File Gateway. For example, if a bucket contains an object named `test` or `test/` and a folder named `test/test1`, only `test/` and `test/test1` are visible in a File Gateway.

- Unless you configure your file share connection to save your user credentials and persist across system restarts, you might need to remount your file share each time you restart your client system.

To mount an SMB file share using Windows File Explorer

1. Press the Windows key and type **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**, type `\\[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 Microsoft AD logon or guest account user password.
7. Choose **Finish** to complete your mount point.

Note

Any files or directories that start with a dot (.) character will be marked hidden in Windows. To make these files and directories visible, you must select the **Hidden items** checkbox on the **View** tab in Windows File Explorer.

You can edit file share settings, edit allowed and denied users and groups, and change the guest access password from the Storage Gateway Management Console. You can also refresh the data in the file share's cache and delete a file share from the console.

To modify your SMB file share's properties

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.

2. On the navigation pane, choose **File Shares**.
3. On the **File Share** page, select the check box by the SMB file share that you want to modify.
4. For Actions, choose the action that you want:
 - Choose **Edit file share settings** to modify share access.
 - Choose **Edit allowed/denied users** to add or delete users and groups, and then type the allowed and denied users and groups into the **Allowed Users**, **Denied Users**, **Allowed Groups**, and **Denied Groups** boxes. Use the **Add Entry** buttons to create new access rights, and the **(X)** button to remove access.

 **Note**

Groups must be prefixed with the @ character. Acceptable formats include: DOMAIN \User1, user1, @group1, and @DOMAIN\group1.

5. When you're finished, choose **Save**.

When you enter allowed users and groups, you are creating an allow list. Without an allow list, all authenticated Microsoft AD users can access the SMB file share. Any users and groups that are marked as denied are added to a deny list and can't access the SMB file share. In instances where a user or group is on both the deny list and allow list, the deny list takes precedence.

You can turn on Access Control Lists(ACLs) on your SMB file share. For information about how to turn on ACLs, see [Using Windows ACLs to limit SMB file share access](#).

Next Step

[Test your S3 File Gateway](#)

Using file shares on buckets with pre-existing objects

You can export a file share on an Amazon S3 bucket with objects created outside of the File Gateway using either NFS or SMB. Objects in the bucket that were created outside of the gateway display as files in either the NFS or SMB file system when your file system clients access them. Standard Portable Operating System Interface (POSIX) access and permissions are used in the file share. When you write files back to an Amazon S3 bucket, the files assume the properties and access rights that you give them.

You can upload objects to an S3 bucket at any time. For the file share to display these newly added objects as files, you need to refresh the S3 bucket. For more information, see [the section called "Refreshing Amazon S3 bucket object cache"](#).

Note

We don't recommend having multiple writers for one Amazon S3 bucket. If you do, be sure to read the section "Can I have multiple writers to my Amazon S3 bucket?" in the [Storage Gateway FAQ](#).

To assign metadata defaults to objects accessed using NFS, see Editing Metadata Defaults in [Managing your Amazon S3 File Gateway](#).

For SMB, you can export a share using Microsoft AD or guest access for an Amazon S3 bucket with pre-existing objects. Objects exported through an SMB file share inherits POSIX ownership and permissions from the parent directory right above it. For objects under the root folder, root Access Control Lists (ACL) are inherited. For Root ACL, the owner is `smbguest` and the permissions for files are 666 and the directories are 777. This applies to all forms of authenticated access (Microsoft AD and guest).

Test your S3 File Gateway

Use the following procedure to test your gateway by copying files and folders to your mapped drive and verifying that they appear in your Amazon S3 bucket automatically.

To upload files from your Windows client to Amazon S3

1. On your Windows client, navigate to the drive that you mounted your file share on. The name of your drive is preceded by the name of your S3 bucket.
2. Copy files or a folder to the drive.
3. On the Amazon S3 Management Console, navigate to your mapped bucket. You should see the files and folders that you copied in the Amazon S3 bucket that you specified.

You can see the file share that you created in the **File shares** tab in the AWS Storage Gateway Management Console.

Your NFS or SMB client can write, read, delete, rename, and truncate files.

 **Note**

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

Keep in mind these points about how File Gateways work with S3:

- Reads are served from a read-through cache. In other words, if data isn't available, it's fetched from S3 and added to the cache.
- Writes are sent to S3 through optimized multipart uploads by using a write-back cache.
- Read and writes are optimized so that only the parts that are requested or changed are transferred over the network.
- Deletes remove objects from S3.
- Directories are managed as folder objects in S3, using the same syntax as in the Amazon S3 console. You can rename empty directories.
- Recursive file system operation performance (for example `ls -l`) depends on the number of objects in your bucket.

Managing your Amazon S3 File Gateway

The topics in this section provide information about how to manage your Amazon S3 File Gateway resources. Gateway management includes granting permissions for your gateway to access file shares and Amazon S3 buckets, editing information and settings for gateways and file shares, deleting file shares, refreshing cached objects, and understanding operational status indicators for gateways and file shares.

Topics

- [Edit basic gateway information](#) - Learn how to use the Storage Gateway console to edit basic information for an existing gateway, including the gateway name, time zone, and CloudWatch log group.
- [Granting access and permissions](#) - Learn how use IAM roles to provide your gateway with access permissions for Amazon S3 buckets and Amazon VPC endpoints, prevent certain security issues, and connect file shares to buckets across AWS accounts.
- [Delete a file share](#) - Learn how to delete a file share using the Storage Gateway console.
- [Editing gateway SMB settings](#) - Learn how to edit gateway-level SMB settings that control security strategy, Active Directory authentication, guest access, local group permissions, and file share visibility for the SMB file shares on a gateway.
- [Edit SMB file share settings](#) - Learn how to edit settings to configure name, logging, cache refresh, storage class, file export, and more for an SMB file share.
- [Limit SMB file share access](#) - Learn how to add allowed or denied users or groups to limit access to your SMB file share.
- [Edit NFS file share settings](#) - Learn how to edit settings to configure name, logging, cache refresh, storage class, file export, and more for an NFS file share.
- [Edit NFS file share metadata defaults](#) - Learn how to edit default metadata values that include Unix permissions for files and folders on NFS files shares.
- [Limit NFS file share access](#) - Learn how to to limit access to clients from specific IP addresses or IP ranges for your NFS fileshare.
- [Refreshing Amazon S3 bucket object cache](#) - Learn how to refresh the S3 bucket object cache for a file share and configure a schedule to refresh the cache automatically.
- [Using S3 Object Lock](#) - Learn about how Amazon S3 File Gateway works with the S3 Object Lock feature.

- [File share status](#) - Learn how to view and interpret file share status.
- [Gateway status](#) - Learn how to view and interpret gateway status.
- [Managing bandwidth for your Amazon S3 File Gateway](#) - Learn how to limit the upload throughput from your gateway to AWS to control the amount of network bandwidth the gateway uses.

Edit basic information for an S3 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**.

Granting access and permissions for file shares and buckets

After your S3 File Gateway is activated and running, you can add additional file shares and grant access to Amazon S3 buckets, including buckets in different AWS accounts than your gateways and file shares. The following sections describe how to use IAM roles to provide your gateway with access permissions for Amazon S3 buckets and VPC endpoints, prevent certain security issues, and connect file shares to buckets across AWS accounts.

For information about how to create a new file share, see [Creating a file share](#).

This section contains the following topics, which provide additional information about how to grant access and permissions for file shares and Amazon S3 buckets:

Topics

- [Granting access to an Amazon S3 bucket](#) - Learn how to grant access for your File Gateway to upload files into your Amazon S3 bucket, and to perform actions on any access points or Amazon Virtual Private Cloud (Amazon VPC) endpoints that it uses to connect to the bucket.
- [Cross-service confused deputy prevention](#) - Learn how to prevent a common security issue where an entity that doesn't have permission to perform an action can coerce a more-privileged entity to perform the action.
- [Using a file share for cross-account access](#) - Learn how to grant access for an Amazon Web Services account and users of that account to access resources that belong to another Amazon Web Services account.

Granting access to an Amazon S3 bucket

When you create a file share, your File Gateway requires access to upload files into your Amazon S3 bucket, and to perform actions on any access points or virtual private cloud (VPC) endpoints that it uses to connect to the bucket. To grant this access, your File Gateway assumes an AWS Identity and Access Management (IAM) role that is associated with an IAM policy that grants this access.

The role requires this IAM policy and a security token service trust (STS) relationship for it. The policy determines which actions the role can perform. In addition, your S3 bucket and any

associated access points or VPC endpoints must have an access policy that allows the IAM role to access them.

You can create the role and access policy yourself, or your File Gateway can create them for you. If your File Gateway creates the policy for you, the policy contains a list of S3 actions. For information about roles and permissions, see [Creating a role to delegate permissions to an AWS service](#) in the *IAM User Guide*.

The following example is a trust policy that allows your File Gateway to assume an IAM role.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "",
      "Effect": "Allow",
      "Principal": {
        "Service": "storagegateway.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

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](#)

If you don't want your File Gateway to create a policy on your behalf, you can create your own policy and attach it to your file share. For more information about how to do this, see [Creating a file share](#).

The following example policy allows your File Gateway to perform all the Amazon S3 actions listed in the policy. The first part of the statement allows all the actions listed to be performed on the S3 bucket named `amzn-s3-demo-bucket`. The second part allows the listed actions on all objects in `amzn-s3-demo-bucket`.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "s3:GetAccelerateConfiguration",
        "s3:GetBucketLocation",
        "s3:GetBucketVersioning",
        "s3:ListBucket",
        "s3:ListBucketVersions",
        "s3:ListBucketMultipartUploads"
      ],
      "Resource": "arn:aws:s3:::amzn-s3-demo-bucket",
      "Effect": "Allow"
    },
    {
      "Action": [
        "s3:AbortMultipartUpload",
        "s3:DeleteObject",
        "s3:DeleteObjectVersion",
        "s3:GetObject",
        "s3:GetObjectAcl",
        "s3:GetObjectVersion",
        "s3:ListMultipartUploadParts",
        "s3:PutObject",
        "s3:PutObjectAcl"
      ],
      "Resource": "arn:aws:s3:::amzn-s3-demo-bucket/*",
      "Effect": "Allow"
    }
  ]
}
```

```
}
```

The following example policy is similar to the preceding one, but allows your File Gateway to perform actions required to access a bucket through an access point.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "s3:AbortMultipartUpload",
        "s3:DeleteObject",
        "s3:DeleteObjectVersion",
        "s3:GetObject",
        "s3:GetObjectAcl",
        "s3:GetObjectVersion",
        "s3:ListMultipartUploadParts",
        "s3:PutObject",
        "s3:PutObjectAcl"
      ],
      "Resource": "arn:aws:s3:us-east-1:111122223333:accesspoint/
TestAccessPointName/*",
      "Effect": "Allow"
    }
  ]
}
```

Note

If you need to connect your file share to an S3 bucket through a VPC endpoint, see [Endpoint policies for Amazon S3](#) in the *AWS PrivateLink User Guide*.

Note

For encrypted buckets, the fileshare must use the key in the destination S3 bucket account.

Note

If your File Gateway uses SSE-KMS or DSSE-KMS for encryption, make sure the IAM role associated with the file share includes `kms:Encrypt`, `kms:Decrypt`, `kms:ReEncrypt*`, `kms:GenerateDataKey`, and `kms:DescribeKey` permissions. For more information, see [Using Identity-Based Policies \(IAM Policies\) for Storage Gateway](#).

Cross-service confused deputy prevention

The confused deputy problem is a security issue where an entity that doesn't have permission to perform an action can coerce a more-privileged entity to perform the action. In AWS, cross-service impersonation can result in the confused deputy problem. Cross-service impersonation can occur when one service (the *calling service*) calls another service (the *called service*). The calling service can be manipulated to use its permissions to act on another customer's resources in a way it should not otherwise have permission to access. To prevent this, AWS provides tools that help you protect your data for all services with service principals that have been given access to resources in your account.

We recommend using the [aws:SourceArn](#) and [aws:SourceAccount](#) global condition context keys in resource policies to limit the permissions that AWS Storage Gateway gives another service to the resource. If you use both global condition context keys, the `aws:SourceAccount` value and the account in the `aws:SourceArn` value must use the same account ID when used in the same policy statement.

The value of `aws:SourceArn` must be the ARN of the Storage Gateway with which your file share is associated.

The most effective way to protect against the confused deputy problem is to use the `aws:SourceArn` global condition context key with the full ARN of the resource. If you don't know the full ARN of the resource or if you are specifying multiple resources, use the `aws:SourceArn` global context condition key with wildcards (*) for the unknown portions of the ARN. For example, `arn:aws:servicename:123456789012:*`.

The following example shows how you can use the `aws:SourceArn` and `aws:SourceAccount` global condition context keys in Storage Gateway to prevent the confused deputy problem.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ConfusedDeputyPreventionExamplePolicy",
      "Effect": "Allow",
      "Principal": {
        "Service": "storagegateway.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "444455556666"
        },
        "ArnLike": {
          "aws:SourceArn": "arn:aws:storagegateway:us-
east-1:444455556666:gateway/sgw-123456DA"
        }
      }
    }
  ]
}
```

Using a file share for cross-account access

Cross-account access is when an Amazon Web Services account and users for that account are granted access to resources that belong to another Amazon Web Services account. With File Gateways, you can use a file share in one Amazon Web Services account to access objects in an Amazon S3 bucket that belongs to a different Amazon Web Services account.

To use a file share owned by one Amazon Web Services account to access an S3 bucket in a different Amazon Web Services account

1. Make sure that the S3 bucket owner has granted your Amazon Web Services account access to the S3 bucket that you need to access and the objects in that bucket. For information about how to grant this access, see [Example 2: Bucket owner granting cross-account bucket permissions](#) in the *Amazon Simple Storage Service User Guide*. For a list of the required permissions, see [Granting access to an Amazon S3 bucket](#).

2. Make sure that the IAM role that your file share uses to access the S3 bucket includes permissions for operations such as `s3:GetObjectAcl` and `s3:PutObjectAcl`. In addition, make sure that the IAM role includes a trust policy that allows your account to assume that IAM role. For an example of such a trust policy, see [Granting access to an Amazon S3 bucket](#).

If your file share uses an existing role to access the S3 bucket, you should include permissions for `s3:GetObjectAcl` and `s3:PutObjectAcl` operations. The role also needs a trust policy that allows your account to assume this role. For an example of such a trust policy, see [Granting access to an Amazon S3 bucket](#).

3. Choose **Gateway files accessible to S3 bucket owner** when creating your file share or editing file share settings in the <https://console.aws.amazon.com/storagegateway/home>.

When you have created or updated your file share for cross-account access and mounted the file share on-premises, we highly recommend that you test your setup. You can do this by listing directory contents or writing test files and making sure the files show up as objects in the S3 bucket.

Important

Make sure to set up the policies correctly to grant cross-account access to the account used by your file share. If you don't, updates to files through your on-premises applications don't propagate to the Amazon S3 bucket that you're working with.

Resources

For additional information about access policies and access control lists, see the following:

[Guidelines for using the available access policy options](#) in the *Amazon Simple Storage Service User Guide*

[Access Control List \(ACL\) overview](#) in the *Amazon Simple Storage Service User Guide*

Delete a file share

If you no longer need a file share, you can delete it from the Storage Gateway console. When you delete a file share, the gateway is detached from the Amazon S3 bucket that the file share maps to. However, the S3 bucket and its contents aren't deleted.

If your gateway is uploading data to a S3 bucket when you delete a file share, the delete process doesn't complete until all the data is uploaded. The file share has the **DELETING** status until the data is completely uploaded.

If you don't want to wait for your data to be completely uploaded, see the **To forcibly delete a file share** procedure later in this topic.

To delete a file share

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. Choose **File shares**, then select one or more file shares to delete.
3. For **Actions**, choose **Delete file share**. The confirmation dialog box appears.
4. Verify that you want to delete the specified file shares, then type the word *delete* in the confirmation box and choose **Delete**.

In certain cases, you might not want to wait until all the data written to files on the Network File System (NFS) file share is uploaded before deleting the file share. For example, you might want to intentionally discard data that was written but has not yet been uploaded, or the Amazon S3 bucket that backs the file share might have already been deleted, meaning that uploading the specified data is no longer possible.

In these cases, you can forcibly delete the file share by using the AWS Management Console or the `DeleteFileShare` API operation. This operation stops the data upload process. When it does, the file share enters the **FORCE_DELETING** status. To forcibly delete a file share using the Storage Gateway console, see the procedure following.

To forcibly delete a file share

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. From the **File shares** list page, choose file share that you flagged for deletion in the procedure above to view its details. After a few seconds, a deletion notification message appears on the **Details** tab.
3. In the message that appears on the **Details** tab, verify the ID of the file share that you want to forcibly delete, select the confirmation box, and choose **Force delete now**.

Note

You cannot undo the force delete operation.

When you forcibly delete a file share, pieces of partially-transferred files from multi-part uploads might remain on Amazon S3 where they can incur storage charges. We recommend configuring an Amazon S3 bucket lifecycle rule to delete these file parts automatically. For more information, see [Best practices: managing multipart uploads](#).

You can also use the [DeleteFileShare](#) API operation to forcibly delete the file share. Deleting a file share using the API requires the `storagegateway:DeleteFileShare` IAM policy permission.

Editing SMB settings for a gateway

Gateway-level SMB settings let you configure the security strategy, Active Directory authentication, guest access, local group permissions, and file share visibility for the SMB file shares on a gateway.

To edit gateway level SMB settings

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 the settings you want to edit.

This section contains the following topics, which provide additional information and procedures related to configuring each of the individual SMB settings for your gateway.

Topics

- [Set gateway security level](#) - Learn how to set a security level to specify connection requirements such as Server Message Block (SMB) signing and encryption, and whether to allow connections from SMB version 1 clients.
- [Configure Active Directory authentication](#) - Learn how to configure your corporate Active Directory or AWS Managed Microsoft AD for user authenticated access to your SMB file share.

- [Provide guest access](#) - Learn how to configure your gateway to allow guest access for any user that provides the correct guest account username and password.
- [Configure local groups](#) - Learn how to configure local groups to grant Active Directory users special file share permissions.
- [Set file share visibility](#) - Learn how to specify whether the shares on a gateway are visible when listing shares to users.

Set a security level for your gateway

By using a S3 File Gateway, you can specify a security level for your gateway. By specifying this security level, you can set whether your gateway should require Server Message Block (SMB) signing or SMB encryption, or whether you want to allow SMB version 1.

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 strategy level can affect performance of the gateway.

- **Enforce AES256 encryption** – If you choose this option, S3 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 all current SMB clients on Microsoft Windows.
- **Enforce encryption** – If you choose this option, S3 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 all current SMB clients on Microsoft Windows.

- **Enforce signing** – If you choose this option, S3 File Gateway only allows connections from SMBv2 or SMBv3 clients that have signing turned on. This option works with all current SMB clients on Microsoft Windows.
- **Client negotiated** – If you choose this option, requests are established based on what is negotiated by the client. This option is recommended when you want to maximize compatibility across different clients in your environment.

 **Note**

For gateways activated before June 20, 2019, the default security level is **Client negotiated**.

For gateways activated on June 20, 2019 and later, the default security level is **Enforce encryption**.

5. Choose **Save**.

Use Active Directory to authenticate users

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

 **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](#).

You can also activate access control lists (ACLs) on your SMB file share. For information about how to activate ACLs, see [Using Windows ACLs to limit SMB file share access](#).

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**.

To limit file share access to specific AD users and groups

1. In the Storage Gateway console, choose the file share that you want to limit access to.
2. From the **Actions** drop-down menu, choose **Edit file share access settings**.
3. In the **User and group file share access** section, choose your settings.

For **Allowed users and groups**, choose **Add allowed user** or **Add allowed group** and enter an AD user or group that you want to allow file share access. Repeat this process to allow as many users and groups as necessary.

For **Denied users and groups**, choose **Add denied user** or **Add denied group** and enter an AD user or group that you want to deny file share access. Repeat this process to deny as many users and groups as necessary.

Note

The **User and group file share access** section appears only if **Active Directory** is selected.

Groups must be prefixed with the @ character. Acceptable formats include: DOMAIN \User1, user1, @group1, and @DOMAIN\group1.

If you configure **Allowed and Denied Users and Groups** lists, then Windows ACLs will not grant any access that overrides those lists.

The **Allowed and Denied Users and Groups** lists are evaluated before ACLs, and control which users can mount or access the file share. If any users or groups are placed on the **Allowed** list, the list is considered active, and only those users can mount the file share.

After a user has mounted a file share, ACLs then provide more granular protection that controls which specific files or folders the user can access. For more information, see [Activating Windows ACLs on a new SMB file share](#).

4. When you finish adding your entries, choose **Save**.

Provide guest access to your file share

You can configure your S3 File Gateway to allow guest access for any user that is able to provide the correct guest account username and password. If you want this to be the only method by which users can access your file gateway, then you do not need to join the gateway to a Microsoft Active Directory domain. You can also use this guest access method to create file shares on an S3 File Gateway that is a member of an Active Directory domain.

When you configure a file share to use the **Guest Access** authentication method, the guest access username is `smbguest`. Before you can create a file share using guest access, you need to change the default password for the `smbguest` user.

You can use the following procedure to change the password for the guest user `smbguest`.

To change the guest access password

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. Choose **Gateways** from the navigation pane on the left side of the console page, and then choose the **Name** of the gateway for which you want to provide guest access.
3. From the **Actions** drop down menu, choose **Edit SMB settings**, and then choose **Guest access settings**.
4. For **Guest password**, enter the guest access password you want to set, and then choose **Save changes**.

Configure local groups for your gateway

Local Group settings allow you to grant Active Directory users or groups special permissions for the SMB file shares on your gateway.

You can use Local Group settings to assign Gateway Admin permissions. Gateway Admins can use the Shared Folders Microsoft Management Console snap-in to force-close files that are open and locked.

Note

You must add at least one Gateway Admin user or group before you can join your gateway to an Active Directory domain.

To assign Gateway Admins

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 **Local Group settings**.
4. In the **Local Group settings** section, choose your settings. This section appears only for file shares that use Active Directory.

For **Gateway Admins**, add Active Directory users and groups that you want to grant local Gateway Admin permissions. Add one user or group per line, including the domain name. For example, **corp\Domain Admins**. To create additional lines, choose **Add new Gateway Admin**.

Note

Editing Gateway Admins disconnects and reconnects all SMB file shares.

5. Choose **Save changes**, then choose **Proceed** to acknowledge the warning message that appears.

Set file share visibility

File share visibility controls whether the shares on a gateway are visible when listing shares to users, such as in a net view or browse list. If the file shares on a gateway are visible, then clients can easily discover the shares using a file browser if they know the gateway IP address or DNS name. If the file shares are not visible, then clients need to know the file share name in addition to the gateway IP or DNS name to be able to discover the shares.

Note

This setting is not an effective method for securing access to the file shares in your deployment. For security, we recommend configuring permissions to limit access to specific users and groups. For instructions, see [Limit user and group access for your SMB file share](#).

To set file share visibility

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 **File share visibility settings**.
4. For **Visibility status**, select the check box if you want the shares on this gateway to appear when the gateway lists shares to users. Keep the check box cleared if you do not want the shares on this gateway to appear when the gateway lists shares to users.

Edit settings for your SMB file share

You can edit the following settings for an existing SMB file share:

- **File share name** - choose a name for the file share
- **Audit logs** - turn audit logs on or off
- **Existing log group list** - choose an existing log group for audit logs
- **Non-gateway file cache refresh time** - specify the interval at which to refresh the file share's cache

Note

Setting this value shorter than 30 minutes can negatively impact gateway performance in situations where large numbers of Amazon S3 objects are frequently created or deleted.

- **Upload events settling time** - specify the number of seconds to wait after the last point in time that a client wrote to a file before generating an ObjectUploaded notification
- **Storage class for new objects** - choose a storage class to use for new objects created in your Amazon S3 bucket
- **Guess MIME type** - choose whether you want Storage Gateway to guess the MIME type for uploaded objects based on file extensions
- **Gateway files accessible to S3 bucket owner** - choose whether to make files on the gateway accessible to the AWS account that owns the Amazon S3 bucket that is linked to the file share

- **Enable requester pays** - choose whether to require accounts that read or request data from the file share to pay for access charges, rather than the bucket owner
- **Export as** - choose whether files are exported in read-write or read-only state
- **File and directory access controlled by** - choose whether to use Windows ACLs or POSIX permissions to control file and directory access
- **Opportunistic lock (oplock)** - choose whether allow the file share to use opportunistic locking to optimize the file buffering strategy
- **Force case sensitivity** - choose whether the client or the gateway controls case sensitivity for file and directory names
- **Access based enumeration for files and directories** - choose whether to make the files and folders on the share visible to all users during directory enumeration, or only to users who have read access

Note

You cannot edit an existing file share to point to a new bucket or access point, or modify the VPC endpoint settings. You can configure those settings only when creating a new file share.

To edit the file share settings

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. Choose **File shares**, and then choose the file share that you want to update.
3. For **Actions**, choose **Edit file share settings**.
4. Edit any settings that you want to change.
5. Choose **Save**.

Limit user and group access for your SMB file share

We recommend adding allowed or denied users or groups to limit access to your file share. If you don't, the file share will be available to all authenticated users.

To edit SMB access settings

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. Choose **File shares**, and then choose the SMB file share that you want to edit.
3. For **Actions**, choose **Edit file share access settings**.
4. In the **User and group file share access** section, choose your settings.

For **Allowed users and groups**, choose **Add allowed user** or **Add allowed group** and enter an AD user or group that you want to allow file share access. Repeat this process to allow as many users and groups as necessary. Any users not in the **Allowed user and groups** list will be denied access.

For **Denied users and groups**, choose **Add denied user** or **Add denied group** and enter an AD user or group that you want to deny file share access. Repeat this process to deny as many users and groups as necessary. If the **Allowed users and groups** list is empty, all users other than those on the **Denied users and groups** list will be granted access.

Note

Enter only the AD user or group name. The domain name is implied by the membership of the gateway in the specific AD that the gateway is joined to. If you don't specify any allowed or denied users or groups, any authenticated AD user can export the file share.

Change the server-side encryption method for an existing file share

The following procedure describes how to change the server-side encryption method for an existing NFS or SMB file share using the Storage Gateway console. To perform this action using the Storage Gateway API, see [UpdateNFSFileShare](#) or [UpdateSMBFileShare](#) in the *AWS Storage Gateway API Reference*.

Note

Updating the encryption method applies the new method to existing objects stored in the Amazon S3 buckets after the update.

If you configure your File Gateway to use SSE-KMS for encryption, you must manually add `kms:Encrypt`, `kms:Decrypt`, `kms:ReEncrypt*`, `kms:GenerateDataKey`, and `kms:DescribeKey` permissions to the IAM role associated with the file share. For more information, see [Using Identity-Based Policies \(IAM Policies\) for Storage Gateway](#).

To change the server-side encryption method for an NFS or SMB file share

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. Choose **File shares**, and then choose the file share for which you want to change the encryption method.
3. For **Actions**, choose **Edit file share encryption**.
4. For **Encryption**, choose the type of encryption you want to use for files at rest in Amazon S3:
 - To use server-side encryption managed with Amazon S3 (SSE-S3), choose **S3-Managed Keys (SSE-S3)**. For more information, see [Using server-side encryption with Amazon S3 managed keys](#) in the *Amazon Simple Storage Service User Guide*.
 - To use server-side encryption managed with AWS Key Management Service (SSE-KMS), choose **KMS-Managed Keys (SSE-KMS)**. For **Primary KMS key**, choose an existing AWS KMS key, or choose **Create a new KMS key** to create a new KMS key in the AWS Key Management Service (AWS KMS) console.

For more information about AWS KMS, see [What is AWS Key Management Service?](#) in the *AWS Key Management Service Developer Guide*.

- To use dual-layer server-side encryption managed with AWS Key Management Service (DSSE-KMS), choose **Dual-layer server-side encryption with AWS Key Management Service keys (DSSE-KMS)**. For **Primary KMS key**, choose an existing AWS KMS key, or choose **Create a new KMS key** to create a new KMS key in the AWS Key Management Service (AWS KMS) console.

For more information about DSSE-KMS, see [Using dual-layer server-side encryption with AWS KMS keys](#) in the *Amazon Simple Storage Service User Guide*.

Note

There are additional charges for using DSSE-KMS and AWS KMS keys. For more information, see [AWS KMS pricing](#).

To specify an AWS KMS key with an alias that is not listed or to use an AWS KMS key from a different AWS account, you must use the AWS Command Line Interface. Asymmetric KMS keys are not supported. For more information, see [CreateSMBFileShare](#) in the *AWS Storage Gateway API Reference*.

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

Edit settings for your NFS file share

Use the following procedure to edit settings for an existing NFS file share after you create it.

Note

You cannot edit an existing file share to point to a new bucket or access point, or modify the VPC endpoint settings. You can configure those settings only when creating a new file share.

To edit the file share settings

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. Choose **File shares**, and then choose the file share that you want to update.
3. For **Actions**, choose **Edit file share settings**.
4. For **File share name**, enter a name for the file share.
5. For **Audit logs**, select one of the following:
 - To create a new log group for this file share, choose **Create a new log group**.
 - To send health and resource notifications for this file share to an existing log group, choose **Use and existing log** group, and then choose the desired group from the list.
 - To turn off logging for this file share, choose **Deactivate logging**.

For more information about audit logs, see [Understanding S3 File Gateway audit logs](#).

6. For **Non-gateway file cache refresh time**, choose **Set refresh interval**, and then set the time in **Minutes** or **Days** to refresh the file share's cache using Time To Live (TTL). TTL is the length of time since the last refresh. After the TTL interval has elapsed, accessing a directory causes the File Gateway to refresh that directory's contents from the Amazon S3 bucket.

 **Note**

Setting this value shorter than 30 minutes can negatively impact gateway performance in situations where large numbers of Amazon S3 objects are frequently created or deleted.

7. For **Upload events settling time**, choose **Set settling time**, and then enter the settling time in seconds. **Settling time** controls the minimum delay between the most recent client write operation and generation of the ObjectUploaded log notification. Because clients can make many small writes to files in a short time, we recommend setting this parameter for as long as possible to avoid generating multiple notifications for the same file in rapid succession. For more information, see [Getting file upload notification](#).
8. For **Storage class for new objects**, choose a storage class from the dropdown list. For more information about storage classes, see [Using storage classes with a File Gateway](#).
9. Under **Object metadata**, do the following:
 - a. Select **Guess MIME type** if you want to allow Storage Gateway to guess the media type for uploaded objects based on their file extensions.
 - b. Select **Gateway files accessible to S3 bucket owner**, if you want the AWS account that owns the S3 bucket to have full ownership of files created by the gateway, including read, write, edit, and delete permissions.
10. Select **Enable requester pays** if you want the file requester rather than the bucket owner to pay the cost of the data request and download from the S3 bucket.
- 11.
12. For **Access level**, choose one of the following:
 - **Root squash (default)**: Access for the remote superuser (root) is mapped to UID (65534) and GID (65534).
 - **All squash**: All user access is mapped to User ID (UID) (65534) and Group ID (GID) (65534).

- **No root squash:** The remote superuser (root) receives access as root.

13. For **Export as**, select one of the following:

- To allow clients to read and write files on the file share, select **Read/Write**.
- To allow clients to read files but not write to the file share, select **Read-only**.

 **Note**

For file shares that are mounted on a Microsoft Windows client, if you choose **Read-only**, you might see a message about an unexpected error keeping you from creating the folder. You can ignore this message.

14. When you are done editing settings, choose **Save changes**.

Edit metadata defaults for your NFS file share

If you don't set metadata values for your files or directories in your bucket, your S3 File Gateway sets default metadata values. These values include Unix permissions for files and folders. You can edit the metadata defaults on the Storage Gateway console.

When your S3 File Gateway stores files and folders in Amazon S3, the Unix file permissions are stored in object metadata. When your S3 File Gateway discovers objects that weren't stored by the S3 File Gateway, these objects are assigned default Unix file permissions. You can find the default Unix permissions in the following table.

Metadata	Description
Directory permissions	The Unix directory mode in the form "nnnn". For example, "0666" represents the access mode for all directories inside the file share. The default value is 0777.
File permissions	The Unix file mode in the form "nnnn". For example, "0666" represents the file mode inside the file share. The default value is 0666.

Metadata	Description
User ID	The default owner ID for files in the file share. The default value is 65534.
Group ID	The default group ID for the file share. The default value is 65534.

To edit metadata defaults

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. Choose **File shares**, and then choose the file share that you want to update.
3. For **Actions**, choose **Edit file metadata defaults**.
4. In the **Edit file metadata defaults** dialog box, provide the metadata information and choose **Save**.

Limit client access for your NFS file share

We recommend editing the NFS client access settings to define a list of specific client IP addresses or CIDR block ranges for NFS clients that are allowed to connect to your NFS file share. If you choose not to limit access, any client on your network can mount to your file share.

To limit client access for your NFS file share

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. Choose **File shares** from the navigation pane on the left side of the console page, and then choose the **File share ID** of the NFS file share that you want to edit.
3. From the **Actions** drop down menu, choose **Edit file share access settings**.

The **Access object** section displays a list of IP addresses and CIDR blocks that are currently allowed to connect to the NFS file share. If access is not currently limited, you will see an entry under **Allowed clients** for the 0.0.0.0/0 CIDR block, which indicates that all possible IPv4 addresses are allowed to connect.

4. Under **Allowed clients**, to the right of the 0.0.0.0/0 CIDR block, choose **Remove**.

5. Choose **Add client**, and then provide an IP address or address range in CIDR notation for the clients that you want to allow.
6. Repeat the previous step to add more IP addresses or ranges as necessary. If make a mistake or need to revoke access, you can choose **Remove** to the right of the IP address or range that you want to delete from the list.
7. Choose **Save changes** when finished.

Refreshing Amazon S3 bucket object cache

As your NFS or SMB client performs file system operations, your gateway maintains an inventory of the objects in the Amazon S3 object cache associated with your file share. Your gateway uses this cached inventory to reduce the latency and frequency of Amazon S3 requests. This operation does not import files into the S3 File Gateway cache storage. It only updates the cached inventory to reflect changes in the inventory of the objects in the Amazon S3 object cache.

To refresh the S3 bucket object cache for your file share, select the method that best fits your use case from the following list, then complete the corresponding procedure below.

Note

Regardless of the method you use, listing a directory for the first time initializes it, which causes the gateway to list the directory's meta data contents from Amazon S3. The time required to initialize a directory is proportional to the number of entries in that directory.

Topics

- [Configure an automated cache refresh schedule using the Storage Gateway console](#)
- [Configure an automated cache refresh schedule using AWS Lambda with an Amazon CloudWatch rule](#)
- [Perform a manual cache refresh using the Storage Gateway console](#)
- [Perform a manual cache refresh using the Storage Gateway API](#)

Configure an automated cache refresh schedule using the Storage Gateway console

The following procedure configures an automatic cache refresh schedule based on a Time To Live (TTL) value that you specify. Before you configure a TTL-based cache refresh schedule, consider the following:

- TTL is measured as the length of time since the last cache refresh for a given directory.
- TTL-based cache refresh occurs only when a given directory is accessed after the specified TTL period has expired.
- The refresh is non-recursive. It occurs only on the specific directories being accessed.
- The refresh incurs Amazon S3 API costs only on directories that have not been synchronized since TTL expiration.
 - Directories are only synchronized if they are accessed by NFS or SMB activity.
 - Synchronization does not occur more frequently than the TTL period that you specify.
- Configuring TTL-based cache refresh is recommended only if you frequently update the contents of your Amazon S3 bucket directly, outside of the workflow between the gateway and the Amazon S3 bucket.
- NFS and SMB operations that access directories with expired TTLs will be blocked while the gateway refreshes the contents of the directory.

Note

Because cache refresh can block directory access operations, we recommend configuring the longest TTL period that is practical for your deployment.

To configure an automated cache refresh schedule using the Storage Gateway console

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. Choose **File shares**.
3. Choose the file share for which you want to configure the refresh schedule.
4. For **Actions**, choose **Edit file share settings**.

5. For **Automated cache refresh from S3 after**, select the check box and set the time in days, hours, and minutes to refresh the file share's cache using Time To Live (TTL). TTL is the length of time since the last refresh after which access to the directory would cause the File Gateway to first refresh that directory's contents from the Amazon S3 bucket.
6. Choose **Save changes**.

Configure an automated cache refresh schedule using AWS Lambda with an Amazon CloudWatch rule

To configure an automated cache refresh schedule using AWS Lambda with an Amazon CloudWatch rule

1. Identify the S3 bucket used by the S3 File Gateway.
2. Check that the *Event* section is blank. It populates automatically later.
3. Create an IAM role, and allow Trust Relationship for Lambda `lambda.amazonaws.com`.
4. Use the following policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "StorageGatewayPermissions",
      "Effect": "Allow",
      "Action": "storagegateway:RefreshCache",
      "Resource": "*"
    },
    {
      "Sid": "CloudWatchLogsPermissions",
      "Effect": "Allow",
      "Action": [
        "logs:CreateLogStream",
        "logs:CreateLogGroup",
        "logs:PutLogEvents"
      ],
      "Resource": "*"
    }
  ]
}
```

```
]
}
```

5. Create a Lambda function from the Lambda console.
6. Use the following function for your Lambda task.

```
import json
import boto3
client = boto3.client('storagegateway')
def lambda_handler(event, context):
    print(event)
    response = client.refresh_cache(
        FileShareARN='arn:aws:storagegateway:ap-southeast-2:672406474878:share/
share-E51FBS9C'
    )
    print(response)
    return 'Your FileShare cache has been refreshed'
```

7. For **Execution role**, choose the IAM role you created.
8. Optional: add a trigger for Amazon S3 and select the event **ObjectCreated** or **ObjectRemoved**.

Note

RefreshCache needs to complete one process before starting another. When you create or delete many objects in a bucket, performance might degrade. Therefore, we recommend against using S3 triggers. Instead, use the Amazon CloudWatch rule described following.

9. Create a CloudWatch rule on the CloudWatch console and add a schedule. Generally, we recommend a *fixed rate* of 30 minutes. However, you can use 1–2 hours on large S3 bucket.
10. Add a new trigger for CloudWatch events and choose the rule you just created.
11. Save your Lambda configuration. Choose **Test**.
12. Choose **S3 PUT** and customize the test to your requirements.
13. The test should succeed. If not, modify the JSON to your requirements and retest.
14. Open the Amazon S3 console, and verify that the event you created and the Lambda function ARN are present.
15. Upload an object to your S3 bucket using the Amazon S3 console or the AWS CLI.

The CloudWatch console generates a CloudWatch output similar to the following.

```
{
  u'Records': [
    {u'eventVersion': u'2.0', u'eventTime': u'2018-09-10T01:03:59.217Z',
    u'requestParameters': {u'sourceIPAddress': u'MY-IP-ADDRESS'},
      u's3': {u'configurationId': u'95a51e1c-999f-485a-b994-9f830f84769f',
    u'object': {u'sequencer': u'00549CC2BF34D47AED', u'key': u'new/filename.jpeg'},
      u'bucket': {u'arn': u'arn:aws:s3:::amzn-s3-demo-bucket', u'name':
    u'MY-GATEWAY-NAME', u'ownerIdentity': {u'principalId': u'A30KNBZ72HVPP9'}}},
    u's3SchemaVersion': u'1.0'},
    u'reponseElements': {u'x-amz-id-2':
    u'76tiugjhvjfyriugiug87t890nefevbck0iA3rPU9I/s4NY9uXwtRL75tCyxasgsdgfsq+IhvAg5M=',
    u'x-amz-request-id': u'651C2D4101D31593'},
      u'awsRegion': u'MY-REGION', u'eventName': u'ObjectCreated:PUT',
    u'userIdentity': {u'principalId': u'AWS:AROAI5LQR5JHFHDFHDFHJ:MY-USERNAME'},
    u'eventSource': u'aws:s3'}
  ]
}
```

The Lambda invocation gives you output similar to the following.

```
{
  u'FileShareARN': u'arn:aws:storagegateway:REGION:ACCOUNT-ID:share/MY-SHARE-
  ID',
  'ResponseMetadata': {'RetryAttempts': 0, 'HTTPStatusCode': 200,
  'RequestId': '6663236a-b495-11e8-946a-bf44f413b71f',
    'HTTPHeaders': {'x-amzn-requestid': '6663236a-b495-11e8-946a-
    bf44f413b71f', 'date': 'Mon, 10 Sep 2018 01:03:59 GMT',
      'content-length': '90', 'content-type': 'application/x-amz-
    json-1.1'
    }
  }
}
```

Your NFS share mounted on your client will reflect this update.

Note

For caches updating large object creation or deletion in large buckets with millions of objects, updates may take hours.

16. Delete your object manually using the Amazon S3 console or AWS CLI.
17. View the NFS share mounted on your client. Verify that your object is gone (because your cache refreshed).
18. Check your CloudWatch logs to see the log of your deletion with the event `ObjectRemoved:Delete`.

```
{
  u'account': u'MY-ACCOUNT-ID', u'region': u'MY-REGION', u'detail': {}, u'detail-
type': u'Scheduled Event', u'source': u'aws.events',
  u'version': u'0', u'time': u'2018-09-10T03:42:06Z', u'id':
u'6468ef77-4db8-0200-82f0-04e16a8c2bdb',
  u'resources': [u'arn:aws:events:REGION:MY-ACCOUNT-ID:rule/FGw-RefreshCache-CW']
}
```

Note

For cron jobs or scheduled tasks, your CloudWatch log event is `u'detail-type': u'Scheduled Event'`.

Perform a manual cache refresh using the Storage Gateway console

To perform a manual cache refresh using the Storage Gateway console

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. Choose **File shares**, and then choose the file share for which you want to perform the refresh.
3. For **Actions**, choose **Refresh cache**.

The time that the refresh process takes depends on the number of objects cached on the gateway and the number of objects that were added to or removed from the S3 bucket.

Perform a manual cache refresh using the Storage Gateway API

The following procedure performs a manual cache refresh using the Storage Gateway API. Before you perform an API-based cache refresh, consider the following:

- You can specify a recursive or non-recursive refresh.
- A recursive refresh is more resource-intensive, and more expensive.
- The refresh incurs Amazon S3 API costs only on directories that you pass as arguments in the request, and descendants of those directories if you specify a recursive refresh.
- The refresh executes concurrently with other operations while the gateway is in use.
 - NFS and SMB operations generally do not become blocked during refreshes, unless a refresh is active for the directory being accessed by the operation.
 - The gateway is unable to determine whether current cache contents are stale, and uses its current contents for NFS and SMB operations regardless of freshness.
 - Because a cache refresh utilizes gateway virtual hardware resources, gateway performance might be negatively impacted while the refresh is in progress.
- Performing API-based cache refresh is recommended only if you update the contents of your Amazon S3 bucket directly, outside of the workflow between the gateway and the Amazon S3 bucket.

Note

If you know the specific directories where you are updating Amazon S3 content outside of the gateway workflow, we recommend specifying these directories in your API-based refresh request to reduce Amazon S3 API costs and gateway performance impact.

To perform a manual cache refresh using the Storage Gateway API

- Send an HTTP POST request to invoke the `RefreshCache` operation with your desired parameters through the Storage Gateway API. For more information, see [RefreshCache](#) in the *AWS Storage Gateway API Reference*.

Note

Sending the `RefreshCache` request only initiates the cache refresh operation. When the cache refresh completes, it doesn't necessarily mean that the file refresh is

complete. To determine that the file refresh operation is complete before you check for new files on the gateway file share, use the `refresh-complete` notification. To do this, you can subscribe to be notified through an Amazon CloudWatch event. For more information, see [Getting notified about file operations](#).

Using S3 Object Lock with Amazon S3 File Gateway

Amazon S3 File Gateway supports accessing S3 buckets that have Amazon S3 Object Lock turned on. Amazon S3 Object Lock allows you to store objects using a "Write Once Read Many" (WORM) model. When you use Amazon S3 Object Lock, you can prevent an object in your S3 bucket from being deleted or overwritten. Amazon S3 Object Lock works together with object versioning to protect your data.

If you turn on Amazon S3 Object Lock, you can still modify the object. For example, it can be written to, deleted, or renamed through a file share on a S3 File Gateway. When you modify an object in this way, S3 File Gateway places a new version of the object without affecting the previous version (that is, the locked object).

For example, If you use the S3 File Gateway NFS or SMB interface to delete a file and the corresponding S3 object is locked, the gateway places an S3 delete marker as the next version of the object, and leaves the original object version in place. Similarly, If a S3 File Gateway modifies the contents or metadata of a locked object, a new version of the object is uploaded with the changes, but the original locked version of the object remains unchanged.

For more information about Amazon S3 Object Lock, see [Locking objects using S3 Object Lock](#) in the *Amazon Simple Storage Service User Guide*.

Understanding file share status

You can view the health of a file share at a glance by looking at its status. If the status indicates that the file share 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 share's status on the Storage Gateway console in the **Status** column. A file share 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 share is configured properly and is available to use. This is the standard status for a file share that's working properly.
CREATING	The file share is not yet fully created and is not ready for use. The CREATING status is transitional. No action is required. If the file share gets stuck in this status, it's probably because the gateway VM lost connection to AWS.
UPDATING	The file share configuration is currently updating. The UPDATING status is transitional. No action is required. If a file share gets stuck in this status, it's probably because the gateway VM lost connection to AWS.
DELETING	The file share is being deleted. The file share is not deleted until all data is uploaded to AWS. The DELETING status is transitional, and no action is required.
FORCE_DELETING	The file share is being deleted forcibly. The file share is deleted immediately and data is not uploaded to AWS. The FORCE_DELETING status is transitional, and no action is required.
UNAVAILABLE	The file share is in an unhealthy state. Action is required. Some possible causes include role policy errors or mapping to an Amazon S3 bucket that doesn't exist. When the issue that caused the unhealthy state is resolved, the file share returns to a status of AVAILABLE.

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	<p>Your gateway might be in an OFFLINE status for one or more of the following reasons:</p> <ul style="list-style-type: none">• 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.

Managing bandwidth for your Amazon S3 File Gateway

You can limit the upload throughput from your gateway to AWS to control the amount of network bandwidth the gateway uses. By default, an activated gateway has no rate limits.

You can configure a bandwidth-rate-limit schedule using the AWS Management Console, an AWS Software Development Kit (SDK), or the AWS Storage Gateway API (see [UpdateBandwidthRateLimitSchedule](#) in the *AWS Storage Gateway API Reference*). Using a bandwidth rate limit schedule, you can configure limits to change automatically throughout the day or week. For more information, see [View and edit the bandwidth-rate-limit schedule for your gateway using the Storage Gateway console](#).

You can monitor your gateway's upload throughput using the `CloudBytesUploaded` metric on the **Monitoring** tab in the Storage Gateway console, or in Amazon CloudWatch.

Note

Bandwidth rate limits apply to Storage Gateway file uploads only. Other gateway operations are not affected.

Bandwidth rate limiting works by balancing the throughput of all files being uploaded, averaged over every second. While it is possible for uploads to cross the bandwidth rate limit briefly for any given micro- or millisecond, this does not typically result in large spikes over longer periods of time.

Configuring bandwidth rate limits and schedules is not currently supported for the Amazon FSx File Gateway type.

Topics

- [View and edit the bandwidth-rate-limit schedule for your gateway using the Storage Gateway console](#)
- [Updating Gateway Bandwidth-Rate Limits Using the AWS SDK for Java](#)
- [Updating Gateway Bandwidth-Rate Limits Using the AWS SDK for .NET](#)
- [Updating Gateway Bandwidth-Rate Limits Using the AWS Tools for Windows PowerShell](#)

View and edit the bandwidth-rate-limit schedule for your gateway using the Storage Gateway console

This section describes how to view and edit the bandwidth rate limit schedule for your gateway.

To view and edit the bandwidth rate limit schedule

1. Open the Storage Gateway console at <https://console.aws.amazon.com/storagegateway/home>.
2. In the left navigation pane, choose **Gateways**, and then choose the gateway that you want to manage.
3. For **Actions**, choose **Edit bandwidth rate limit schedule**.

The gateway's current bandwidth-rate-limit schedule is displayed on the **Edit bandwidth rate limit schedule** page. By default, a new gateway has no defined bandwidth-rate limits.

4. (Optional) Choose **Add new bandwidth rate limit** to add a new configurable interval to the schedule. For each interval you add, enter the following information:
 - **Upload rate** – Enter the upload rate limit, in megabits per second (Mbps). The minimum value is 100 Mbps.

- **Days of week** – Select the day or days during each week when you want the interval to apply. You can apply the interval on weekdays (Monday through Friday), weekends (Saturday and Sunday), every day of the week, or on one specific day each week. To apply the bandwidth-rate limit uniformly and constantly on all days and at all times, choose **No schedule**.
- **Start time** – Enter the start time for the bandwidth interval, using the HH:MM format and the time-zone offset from UTC for your gateway.

 **Note**

Your bandwidth-rate-limit interval begins at the start of the minute that you specify here.

- **End time** – Enter the end time for the bandwidth interval, using the HH:MM format and the time-zone offset from GMT for your gateway.

 **Important**

The bandwidth-rate-limit interval ends at the end of the minute specified here. To schedule an interval that ends at the end of an hour, enter **59**.

To schedule consecutive continuous intervals, transitioning at the start of the hour, with no interruption between the intervals, enter **59** for the end minute of the first interval. Enter **00** for the start minute of the succeeding interval.

5. (Optional) Repeat the previous step as necessary until your bandwidth-rate-limit schedule is complete. If you need to delete an interval from your schedule, choose **Remove**.

 **Important**

Bandwidth-rate-limit intervals cannot overlap. The start time of an interval must occur after the end time of a preceding interval, and before the start time of a following interval.

6. When finished, choose **Save changes**.

Updating Gateway Bandwidth-Rate Limits Using the AWS SDK for Java

By updating bandwidth-rate limits programmatically, you can adjust these limits automatically over a period of time—for example, by using scheduled tasks. The following example demonstrates how to update a gateway's bandwidth-rate limits using the AWS SDK for Java. To use the example code, you should be familiar with running a Java console application. For more information, see [Getting Started](#) in the *AWS SDK for Java Developer Guide*.

Example : Updating Gateway Bandwidth-Rate Limits Using the AWS SDK for Java

The following Java code example updates a gateway's bandwidth-rate limits. To use this example code, you must provide the service endpoint, your gateway Amazon Resource Name (ARN), and the upload limit. For 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*.

```
import java.io.IOException;

import com.amazonaws.AmazonClientException;
import com.amazonaws.auth.PropertiesCredentials;
import com.amazonaws.services.storagegateway.AWSStorageGatewayClient;
import com.amazonaws.services.storagegateway.model.
UpdateBandwidthRateLimitScheduleRequest;
import com.amazonaws.services.storagegateway.model.
UpdateBandwidthRateLimitScheduleReturn;

import java.util.Arrays;
import java.util.Collections;
import java.util.List;

public class UpdateBandwidthExample {

    public static AWSStorageGatewayClient sgClient;

    // The gatewayARN
    public static String gatewayARN = "**** provide gateway ARN ****";

    // The endpoint
    static String serviceURL = "https://storagegateway.us-east-1.amazonaws.com";

    // Rates
    static long uploadRate = 100 * 1024 * 1024; // Bits per second, minimum 100
Megabits/second
```

```

public static void main(String[] args) throws IOException {

    // Create a Storage Gateway client
    sgClient = new AWSStorageGatewayClient(new PropertiesCredentials(
UpdateBandwidthExample.class.getResourceAsStream("AwsCredentials.properties")));
    sgClient.setEndpoint(serviceURL);

    UpdateBandwidth(gatewayARN, uploadRate, null); // download rate not
supported by S3 File Gateways

}

private static void UpdateBandwidth(String gatewayArn, long uploadRate, long
downloadRate) {
    try
    {
        BandwidthRateLimit bandwidthRateLimit = new
BandwidthRateLimit(downloadRate, uploadRate);
        BandwidthRateLimitInterval noScheduleInterval = new
BandwidthRateLimitInterval()
            .withBandwidthRateLimit(bandwidthRateLimit)
            .withDaysOfWeek(Arrays.asList(1, 2, 3, 4, 5, 6, 0))
            .withStartHourOfDay(0)
            .withStartMinuteOfHour(0)
            .withEndHourOfDay(23)
            .withEndMinuteOfHour(59);
        UpdateBandwidthRateLimitScheduleRequest
updateBandwidthRateLimitScheduleRequest =
            new UpdateBandwidthRateLimitScheduleRequest()
            .withGatewayARN(gatewayArn)
            .with
BandwidthRateLimitIntervals(Collections.singletonList(noScheduleInterval));

        UpdateBandwidthRateLimitScheduleReturn
updateBandwidthRateLimitScheuduleResponse =
sgClient.UpdateBandwidthRateLimitSchedule(updateBandwidthRateLimitScheduleRequest);

        String returnGatewayARN =
updateBandwidthRateLimitScheuduleResponse.getGatewayARN();
        System.out.println("Updated the bandwidth rate limits of " +
returnGatewayARN);
    }
}

```

```

        System.out.println("Upload bandwidth limit = " + uploadRate + " bits
per second");
    }
    catch (AmazonClientException ex)
    {
        System.err.println("Error updating gateway bandwidth.\n" +
ex.toString());
    }
}
}

```

Updating Gateway Bandwidth-Rate Limits Using the AWS SDK for .NET

By updating bandwidth-rate limits programmatically, you can adjust these limits automatically over a period of time—for example, by using scheduled tasks. The following example demonstrates how to update a gateway's bandwidth-rate limits by using the AWS Software Development Kit (SDK) for .NET. To use the example code, you should be familiar with running a .NET console application. For more information, see [Getting Started](#) in the *AWS SDK for .NET Developer Guide*.

Example : Updating Gateway Bandwidth-Rate Limits by Using the AWS SDK for .NET

The following C# code example updates a gateway's bandwidth-rate limits. To use this example code, you must provide the service endpoint, your gateway Amazon Resource Name (ARN), and the upload limit. For 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*.

```

using System;
using System.Collections.Generic;
using System.Linq;
using System.Text;
using Amazon.StorageGateway;
using Amazon.StorageGateway.Model;

namespace AWSStorageGateway
{
    class UpdateBandwidthExample
    {
        static AmazonStorageGatewayClient sgClient;
        static AmazonStorageGatewayConfig sgConfig;

        // The gatewayARN
        public static String gatewayARN = "**** provide gateway ARN ****";
    }
}

```

```

        // The endpoint
        static String serviceURL = "https://storagegateway.us-
east-1.amazonaws.com";

        // Rates
        static long uploadRate = 100 * 1024 * 1024; // Bits per second, minimum
        100 Megabits/second

        public static void Main(string[] args)
        {
            // Create a Storage Gateway client
            sgConfig = new AmazonStorageGatewayConfig();
            sgConfig.ServiceURL = serviceURL;
            sgClient = new AmazonStorageGatewayClient(sgConfig);

            UpdateBandwidth(gatewayARN, uploadRate, null);

            Console.WriteLine("\nTo continue, press Enter.");
            Console.Read();
        }

        public static void UpdateBandwidth(string gatewayARN, long uploadRate, long
downloadRate)
        {
            try
            {
                BandwidthRateLimit bandwidthRateLimit = new
BandwidthRateLimit(downloadRate, uploadRate);
                BandwidthRateLimitInterval noScheduleInterval = new
BandwidthRateLimitInterval()
                    .withBandwidthRateLimit(bandwidthRateLimit)
                    .withDaysOfWeek(Arrays.asList(1, 2, 3, 4, 5, 6, 0))
                    .withStartHourOfDay(0)
                    .withStartMinuteOfHour(0)
                    .withEndHourOfDay(23)
                    .withEndMinuteOfHour(59);
                List <BandwidthRateLimitInterval> bandwidthRateLimitIntervals = new
List<BandwidthRateLimitInterval>();
                bandwidthRateLimitIntervals.Add(noScheduleInterval);
                UpdateBandwidthRateLimitScheduleRequest
updateBandwidthRateLimitScheduleRequest =
                    new UpdateBandwidthRateLimitScheduleRequest()
                        .withGatewayARN(gatewayARN)

```

```

        .with BandwidthRateLimitIntervals(bandwidthRateLimitIntervals);

        UpdateBandwidthRateLimitScheduleReturn
updateBandwidthRateLimitScheduleResponse =
sgClient.UpdateBandwidthRateLimitSchedule(updateBandwidthRateLimitScheduleRequest);
        String returnGatewayARN =
updateBandwidthRateLimitScheduleResponse.GatewayARN;
        Console.WriteLine("Updated the bandwidth rate limits of " +
returnGatewayARN);
        Console.WriteLine("Upload bandwidth limit = " + uploadRate + " bits
per second");
    }
    catch (AmazonStorageGatewayException ex)
    {
        Console.WriteLine("Error updating gateway bandwidth.\n" +
ex.ToString());
    }
}
}
}

```

Updating Gateway Bandwidth-Rate Limits Using the AWS Tools for Windows PowerShell

By updating bandwidth-rate limits programmatically, you can adjust these limits automatically over a period of time—for example, by using scheduled tasks. The following example demonstrates how to update a gateway's bandwidth-rate limits using the AWS Tools for Windows PowerShell. To use the example code, you should be familiar with running a PowerShell script. For more information, see [Getting Started](#) in the *AWS Tools for PowerShell User Guide*.

Example : Updating Gateway Bandwidth-Rate Limits by Using the AWS Tools for Windows PowerShell

The following PowerShell script example updates a gateway's bandwidth-rate limits. To use this example script, you must provide your gateway Amazon Resource Name (ARN) and the upload limit.

```

<#
.DESCRIPTION
    Update Gateway bandwidth limits schedule

```

.NOTES

PREREQUISITES:

- 1) AWS Tools for PowerShell from <https://aws.amazon.com/powershell/>
- 2) Credentials and region stored in session using Initialize-AWSDefault.

For more info, see <https://docs.aws.amazon.com/powershell/latest/userguide/specifying-your-aws-credentials.html>

.EXAMPLE

```
powershell.exe .\SG_UpdateBandwidth.ps1
```

```
#>
```

```
$UploadBandwidthRate = 100 * 1024 * 1024
```

```
$gatewayARN = "**** provide gateway ARN ****"
```

```
$bandwidthRateLimitInterval = New-Object
```

```
Amazon.StorageGateway.Model.BandwidthRateLimitInterval
```

```
$bandwidthRateLimitInterval.StartHourOfDay = 0
```

```
$bandwidthRateLimitInterval.StartMinuteOfHour = 0
```

```
$bandwidthRateLimitInterval.EndHourOfDay = 23
```

```
$bandwidthRateLimitInterval.EndMinuteOfHour = 59
```

```
$bandwidthRateLimitInterval.DaysOfWeek = 0,1,2,3,4,5,6
```

```
$bandwidthRateLimitInterval.AverageUploadRateLimitInBitsPerSec =
```

```
$UploadBandwidthRate
```

```
#Update Bandwidth Rate Limits
```

```
Update-SGBandwidthRateLimitSchedule -GatewayARN $gatewayARN `
```

```
-BandwidthRateLimitInterval
```

```
@($bandwidthRateLimitInterval)
```

```
$schedule = Get-SGBandwidthRateLimitSchedule -GatewayARN $gatewayARN
```

```
Write-Output("`nGateway: " + $gatewayARN);
```

```
Write-Output("`nNew bandwidth throttle schedule: " +
```

```
$schedule.BandwidthRateLimitIntervals.AverageUploadRateLimitInBitsPerSec)
```

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 S3 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 S3 File Gateway, see

[Understanding gateway metrics](#) and [Understanding file share 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 shares unavailable: `FileSharesUnavailable` ≥ 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 S3 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 enters 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 on 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 S3 File Gateway

You can monitor your S3 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 S3 File Gateway health logs with CloudWatch log groups](#)
- [Using Amazon CloudWatch metrics](#)
- [Getting notified about file operations](#)
- [Understanding gateway metrics](#)
- [Understanding file share metrics](#)
- [Understanding S3 File Gateway audit logs](#)

Getting S3 File Gateway health logs with CloudWatch log groups

You can use Amazon CloudWatch Logs to get information about the health of your S3 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 S3 File Gateway fails to upload files to an Amazon S3 bucket. 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 S3 File Gateway](#). For general information about CloudWatch log groups, see [Working with Log Groups and Log Streams](#) in the *Amazon CloudWatch User Guide*.

The following is an example of an error reported by an S3 File Gateway.

```
{
  "severity": "ERROR",
  "bucket": "bucket-smb-share2",
  "roleArn": "arn:aws:iam::123456789012:role/amzn-s3-demo-bucket",
  "source": "share-E1A2B34C",
  "type": "InaccessibleStorageClass",
  "operation": "S3Upload",
  "key": "myFolder/myFile.text",
  "gateway": "sgw-B1D123D4",
  "timestamp": "1565740862516"
}
```

This error means that the S3 File Gateway is unable to upload the object `myFolder/myFile.text` to Amazon S3 because it has transitioned out of the Amazon S3 Standard storage class to either the S3 Glacier Flexible Retrieval or the S3 Glacier Deep Archive storage class.

In the preceding gateway health log, these items specify the given information:

- `source: share-E1A2B34C` indicates the file share that encountered this error.
- `"type": "InaccessibleStorageClass"` indicates the type of error that occurred. In this case, this error was encountered when the gateway was trying to upload the specified object to Amazon S3 or read from Amazon S3. However, in this case, the object has transitioned to Amazon S3 Glacier. The value of `"type"` can be any error that the S3 File Gateway encounters. For a list of possible errors, see [Troubleshooting: File Gateway issues](#).
- `"operation": "S3Upload"` indicates that this error occurred when the gateway was trying to upload this object to S3.
- `"key": "myFolder/myFile.text"` indicates the object that caused the failure.
- `gateway": "sgw-B1D123D4"` indicates the S3 File Gateway that encountered this error.
- `"timestamp": "1565740862516"` indicates the time that the error occurred.

For information about how to troubleshoot the errors that may be reported by S3 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 S3 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 S3 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	<code>GatewayId</code> , <code>GatewayName</code>	<p>These dimensions filter for metric data that describes aspects of the gateway. You can identify a S3 File Gateway to work with by specifying both the <code>GatewayId</code> and the <code>GatewayName</code> dimensions.</p> <p>Throughput and latency data of a gateway are based on all the file shares in the gateway.</p> <p>Data is available automatically in 5-minute periods at no charge.</p>

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](#)

Getting notified about file operations

Storage Gateway can initiate the following CloudWatch Events when your file operations are done:

- You can get notified when the gateway finishes the asynchronous uploading of your files from the file share to Amazon S3. Use the `NotificationPolicy` parameter to request a file upload notification. This sends a notification for each completed file upload to Amazon S3. For more information, see [Getting file upload notification](#).
- You can get notified when the gateway finishes the asynchronous uploading of your working file set from the file share to Amazon S3. Use the `NotifyWhenUploaded` API operation to request a working file set upload notification. This sends a notification when all files in the working file set have been uploaded to Amazon S3. For more information, see [Getting working file set upload notification](#).
- You can get notified when the gateway finishes refreshing the cache for your S3 bucket. When you invoke the `RefreshCache` operation through the Storage Gateway console or API, subscribe to the notification when the operation is complete. For more information, see [Getting refresh cache notification](#).

When the file operation you requested is done, Storage Gateway sends you a notification through CloudWatch Events. You can configure CloudWatch Events to send the notification through event targets such as Amazon SNS, Amazon SQS, or an AWS Lambda function. For example, you can configure an Amazon SNS target to send the notification to Amazon SNS consumers such as an email or text message. For information about CloudWatch Events, see [What is CloudWatch Events?](#)

To set up CloudWatch Events notification

1. Create a target, such as an Amazon SNS topic or Lambda function, to invoke when the event you requested in Storage Gateway occurs.
2. Create a rule in the CloudWatch Events console to invoke targets based on an event in Storage Gateway.
3. In the rule, create an event pattern for the event type. The notification sent when the event matches this rule pattern.
4. Select the target and configure the settings.

The following example shows a rule that initiates the specified event type in the specified gateway and in the specified AWS Region. For example, you could specify the Storage Gateway File Upload Event as the event type.

```
{
  "source": [
    "aws.storagegateway"
  ],
  "resources": [
    "arn:aws:storagegateway:AWS Region:account-id
      :gateway/gateway-id"
  ],
  "detail-type": [
    "Event type"
  ]
}
```

For information about how to use CloudWatch Events rules, see [Creating a CloudWatch Events rule that triggers on an event](#) in the *Amazon CloudWatch Events User Guide*.

Getting file upload notification

There are two use cases in which you can use file upload notification:

- For automating in-cloud processing of files that are uploaded, you can call the `NotificationPolicy` parameter and get back a notification ID. The notification that occurs when the files have been uploaded has the same notification ID as the one that was returned by the API. If you map this notification ID to track the list of files that you are uploading, you can initiate processing of the file that is uploaded in AWS when the event with the same ID is generated.
- For content distribution use cases, you can have two S3 File Gateways that map to the same Amazon S3 bucket. The file share client for Gateway1 could upload new files to Amazon S3, and the files are read by file share clients on Gateway2. The files upload to Amazon S3, but they are not visible to Gateway2 because it uses a locally cached version of files in Amazon S3. To make the files visible in Gateway2, you can use the `NotificationPolicy` parameter to request file upload notification from Gateway1 to notify you when the upload file is done. You can then use CloudWatch Events to automatically issue a [RefreshCache](#) request for the file share on Gateway2. When the [RefreshCache](#) request is complete, the new file is visible in Gateway2.

Example Example—File upload notification

The following example shows a file upload notification that is sent to you through CloudWatch when the event matches the rule you created. This notification is in JSON format. You can configure this notification to be delivered to the target as a text message. The detail-type is Storage Gateway Object Upload Event.

```
{
  "version": "0",
  "id": "2649b160-d59d-c97f-3f64-8aaa9ea6aed3",
  "detail-type": "Storage Gateway Object Upload Event",
  "source": "aws.storagegateway",
  "account": "123456789012",
  "time": "2020-11-05T12:34:56Z",
  "region": "us-east-1",
  "resources": [
    "arn:aws:storagegateway:us-east-1:123456789011:share/share-F123D451",
    "arn:aws:storagegateway:us-east-1:123456789011:gateway/sgw-712345DA",
    "arn:aws:s3:::do-not-delete-bucket"
  ],
  "detail": {
    "object-size": 1024,
    "modification-time": "2020-01-05T12:30:00Z",
    "object-key": "my-file.txt",
    "event-type": "object-upload-complete",
    "prefix": "prefix/",
    "bucket-name": "amzn-s3-demo-bucket",
  }
}
```

Field names	Description
version	The current version of the IAM policy.
id	The ID that identifies the IAM policy.
detail-type	A description of the event that initiated the notification that was sent.
source	The AWS service that is the source of the request and notification.

Field names	Description
account	The ID of the AWS account where the request and notification were generated from.
time	When the request to upload files to Amazon S3 was made.
region	The AWS Region where the request and notification was sent from.
resources	The Storage Gateway resources that the policy applies to.
object-size	The size of the object in bytes.
modification-time	The time the client modified the file.
object-key	The path to the file.
event-type	The CloudWatch Events that initiated the notification.
prefix	The prefix name of the S3 bucket.
bucket-name	The name of the S3 bucket.

Getting working file set upload notification

There are two use cases in which you can use the working file set upload notification:

- For automating in-cloud processing of files that are uploaded, you can call the `NotifyWhenUploaded` API and get back a notification ID. The notification that occurs when the working set of files have been uploaded has the same notification ID as the one that was returned by the API. If you map this notification ID to track the list of files that you are uploading, you can initiate processing of the working set of files that are uploaded in AWS when the event with the same ID is generated.
- For content distribution use cases, you can have two S3 File Gateways that map to the same Amazon S3 bucket. The file share client for Gateway1 can upload new files to Amazon S3, and

the files are read by file share clients on Gateway2. The files upload to Amazon S3, but they aren't visible to Gateway2 because it uses a locally cached version of files in S3. To make the files visible in Gateway2, use the [NotifyWhenUploaded](#) API operation to request file upload notification from Gateway1, to notify you when the upload of the working set of files is done. You can then use the CloudWatch Events to automatically issue a [RefreshCache](#) request for the file share on Gateway2. When the [RefreshCache](#) request is complete, the new files are visible in Gateway2. This operation does not import files into the gateway cache storage. It only updates the cached inventory to reflect changes in the inventory of the objects in the S3 bucket.

Example Example—Working file set upload notification

The following example shows a working file set upload notification that is sent to you through CloudWatch when the event matches the rule you created. This notification is in JSON format. You can configure this notification to be delivered to the target as a text message. The `detail-type` is Storage Gateway File Upload Event.

```
{
  "version": "2012-10-17",
  "id": "2649b160-d59d-c97f-3f64-8aaa9ea6aed3",
  "detail-type": "Storage Gateway File Upload Event",
  "source": "aws.storagegateway",
  "account": "123456789012",
  "time": "2017-11-06T21:34:42Z",
  "region": "us-east-2",
  "resources": [
    "arn:aws:storagegateway:us-east-2:123456789011:share/share-F123D451",
    "arn:aws:storagegateway:us-east-2:123456789011:gateway/sgw-712345DA"
  ],
  "detail": {
    "event-type": "upload-complete",
    "notification-id": "11b3106b-a18a-4890-9d47-a1a755ef5e47",
    "request-received": "2018-02-06T21:34:42Z",
    "completed": "2018-02-06T21:34:53Z"
  }
}
```

Field names	Description
version	The current version of the IAM policy.

Field names	Description
id	The ID that identifies the IAM policy.
detail-type	A description of the event that initiated the notification that was sent.
source	The AWS service that is the source of the request and notification.
account	The ID of the AWS account where the request and notification were generated from.
time	When the request to upload files to Amazon S3 was made.
region	The AWS Region where the request and notification was sent from.
resources	The Storage Gateway resources that the policy applies to.
event-type	The CloudWatch Events that initiated the notification.
notification-id	The randomly generated ID of the notification that was sent. This ID is in UUID format. This is the notification ID that is returned when <code>NotifyWhenUploaded</code> is called.
request-received	When the gateway received the <code>NotifyWhenUploaded</code> request.
completed	When all the files in the working-set were uploaded to Amazon S3.

Getting refresh cache notification

For refresh cache notification use case, you can have two S3 File Gateways that map to the same Amazon S3 bucket and the NFS client for Gateway1 uploads new files to the S3 bucket. The files upload to Amazon S3, but they don't appear in Gateway2 until you refresh the cache. This is because Gateway2 uses a locally cached version of the files in Amazon S3. You might want to do something with the files in Gateway2 when the refresh cache is done. Large files could take a while to show up in Gateway2, so you might want to be notified when the cache refresh is done. You can request refresh cache notification from Gateway2 to notify you when all the files are visible in Gateway2.

Example Example—Refresh cache notification

The following example shows a refresh cache notification that is sent to you through CloudWatch when the event matches the rule you created. This notification is in JSON format. You can configure this notification to be delivered to the target as a text message. The `detail-type` is Storage Gateway Refresh Cache Event.

```
{
  "version": "2012-10-17",
  "id": "2649b160-d59d-c97f-3f64-8aaa9ea6aed3",
  "detail-type": "Storage Gateway Refresh Cache Event",
  "source": "aws.storagegateway",
  "account": "209870788375",
  "time": "2017-11-06T21:34:42Z",
  "region": "us-east-2",
  "resources": [
    "arn:aws:storagegateway:us-east-2:123456789011:share/share-F123D451",
    "arn:aws:storagegateway:us-east-2:123456789011:gateway/sgw-712345DA"
  ],
  "detail": {
    "event-type": "refresh-complete",
    "notification-id": "1c14106b-a18a-4890-9d47-a1a755ef5e47",
    "started": "2018-02-06T21:34:42Z",
    "completed": "2018-02-06T21:34:53Z",
    "folderList": [
      "/"
    ]
  }
}
```

Field names	Description
version	The current version of the IAM policy.
id	The ID that identifies the IAM policy.
detail-type	A description of the type of the event that initiated the notification that was sent.
source	The AWS service that is the source of the request and notification.
account	The ID of the AWS account where the request and notification were generated from.
time	When the request to refresh the files in working-set was made.
region	The AWS Region where the request and notification was sent from.
resources	The Storage Gateway resources that the policy applies to.
event-type	The CloudWatch Events that initiated the notification.
notification-id	The randomly generated ID of the notification that was sent. This ID is in UUID format. This is the notification ID that is returned when you call <code>RefreshCache</code> .
started	when the gateway received the <code>RefreshCache</code> request and the refresh was started.
completed	When the refresh of the working-set was completed.

Field names	Description
folderList	A comma-separated list of the paths of folders that were refreshed in the cache. The default is ["/"].

Understanding gateway metrics

The following table describes metrics that cover S3 File Gateways. Each gateway has a set of metrics associated with it. Some gateway-specific metrics have the same name as certain file-share-specific metrics. These metrics represent the same kinds of measurements, but are scoped to the gateway rather than the file share.

Always specify whether you want to work with a gateway or a file share 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 S3 File Gateways.

Metric	Description
AuditNotifications	This metric reports the number of audit logs emitted. Units: Count
AvailabilityNotifications	This metric reports the number of availability-related health notifications that were generated by the gateway in the reporting period.

Metric	Description
	Units: Count
CacheFileSize	<p>This metric tracks the size of files in the gateway cache.</p> <p>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.</p> <p>Units: Bytes</p>
CacheFree	<p>This metric reports the number of available bytes in the gateway cache.</p> <p>Units: Bytes</p>
CacheHitPercent	<p>Percent of application read operations from the gateway that are served from cache. The sample is taken at the end of the reporting period.</p> <p>When there are no application read operations from the gateway, this metric reports 100 percent.</p> <p>Units: Percent</p>

Metric	Description
CachePercentDirty	<p>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.</p> <p>Use this metric with the Sum statistic.</p> <p>Ideally, this metric should remain low.</p> <p>Units: Percent</p>
CachePercentUsed	<p>The percent of the data cache used across the entire gateway. The sample is taken at the end of the reporting period.</p> <p>Units: Percent</p>
CacheUsed	<p>This metric reports the number of used bytes in the gateway cache.</p> <p>Units: Bytes</p>
CloudBytesDownloaded	<p>The total number of bytes that the gateway downloaded from AWS during the reporting period.</p> <p>Use this metric with the Sum statistic to measure throughput and with the Samples statistic to measure IOPS.</p> <p>Units: Bytes</p>

Metric	Description
CloudBytesUploaded	<p>The total number of bytes that the gateway uploaded to AWS during the reporting period.</p> <p>Use this metric with the Sum statistic to measure throughput and with the Samples statistic to measure input/output operations per second (IOPS).</p> <p>Units: Bytes</p>
FilesFailingUpload	<p>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.</p> <p>Use this metric with the Sum statistic to show the number of files which are currently failing to upload to AWS.</p> <p>Units: Count</p>
FileSharesUnavailable	<p>This metric provides the number of file shares on this gateways which are in the Unavailable state.</p> <p>If this metric reports any file shares are unavailable, 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.</p> <p>Units: Count</p>

Metric	Description
FilesRenamed	<p>This metric tracks the number of files renamed in the reporting period.</p> <p>Units: Count</p>
HealthNotifications	<p>This metric reports the number of health notifications that were generated by this gateway in the reporting period.</p> <p>Units: Count</p>
IndexEvictions	<p>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.</p> <p>Units: Count</p>
IndexFetches	<p>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.</p> <p>Units: Count</p>
IoWaitPercent	<p>This metric reports the percentage of time that the CPU is waiting for a response from the local disk.</p> <p>Units: Percent</p>
MemTotalBytes	<p>This metric reports the total amount of memory on the gateway.</p> <p>Units: Bytes</p>

Metric	Description
MemUsedBytes	<p>This metric reports the amount of used memory on the gateway.</p> <p>Units: Bytes</p>
NfsSessions	<p>This metric reports the number of NFS sessions that are active on the gateway.</p> <p>Units: Count</p>
RootDiskFreeBytes	<p>This metric reports the number of available bytes on the root disk of the gateway.</p> <p>If this metric reports less than 20 GB are free, you should increase the size of the root disk.</p> <p>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.</p> <p>Units: Bytes</p>
S3GetObjectRequestTime	<p>This metric reports the time for the gateway to complete S3 get object requests.</p> <p>Units: Milliseconds</p>
S3PutObjectRequestTime	<p>This metric reports the time for the gateway to complete S3 put object requests.</p> <p>Units: Milliseconds</p>
S3UploadPartRequestTime	<p>This metric reports the time for the gateway to complete S3 upload part requests.</p> <p>Units: Milliseconds</p>

Metric	Description
SmbV1Sessions	This metric reports the number of SMBv1 sessions that are active on the gateway. Units: Count
SmbV2Sessions	This metric reports the number of SMBv2 sessions that are active on the gateway. Units: Count
SmbV3Sessions	This metric reports the number of SMBv3 sessions that are active on the gateway. Units: Count
TotalCacheSize	This metric reports the total size of the cache. Units: Bytes
UserCpuPercent	This metric reports the percentage of time that is spent on gateway processing. Units: Percent

Understanding file share metrics

You can find information following about the Storage Gateway metrics that cover file shares. Each file share has a set of metrics associated with it. Some file share-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 share instead.

Always specify whether you want to work with either a gateway or a file share metric before working with a metric. Specifically, when working with file share metrics, you must specify the `File share ID` that identifies the file share 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	<p>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.</p> <p>When there are no application read operations from the file share, this metric reports 100 percent.</p> <p>Units: Percent</p>
CachePercentDirty	<p>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.</p> <p>Use this metric with the Sum statistic.</p> <p>Ideally, this metric should remain low.</p> <div>Note Use the <code>CachePercentDirty</code> metric of the gateway to view the overall percentage of the gateway's cache that has not been persisted to AWS.</div>

Metric	Description
	Units: Percent
CachePercentUsed	<p>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.</p> <p>Units: Percent</p>
CloudBytesUploaded	<p>The total number of bytes that the gateway uploaded to AWS during the reporting period.</p> <p>Use this metric with the Sum statistic to measure throughput and with the Samples statistic to measure IOPS.</p> <p>Units: Bytes</p>
CloudBytesDownloaded	<p>The total number of bytes that the gateway downloaded from AWS during the reporting period.</p> <p>Use this metric with the Sum statistic to measure throughput and with the Samples statistic to measure input/output operations per second (IOPS).</p> <p>Units: Bytes</p>

Metric	Description
FilesFailingUpload	<p>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.</p> <p>Use this metric with the Sum statistic to show the number of files which are currently failing to upload to AWS.</p> <p>Units: Count</p>
ReadBytes	<p>The total number of bytes read from your on-premises applications in the reporting period for a file share.</p> <p>Use this metric with the Sum statistic to measure throughput and with the Samples statistic to measure IOPS.</p> <p>Units: Bytes</p>
WriteBytes	<p>The total number of bytes written to your on-premises applications in the reporting period.</p> <p>Use this metric with the Sum statistic to measure throughput and with the Samples statistic to measure IOPS.</p> <p>Units: Bytes</p>

Understanding S3 File Gateway audit logs

Amazon S3 File Gateway (S3 File Gateway) audit logs provide you with details about user access to files and folders within a file share. You can use them to monitor user activities and take action if inappropriate activity patterns are identified.

Operations

The following table describes the S3 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 S3 File Gateway audit log file access attributes.

Attribute	Definition
accessMode	The permission setting for the object.
accountDomain (SMB only)	The Active Directory (AD) domain that the client's account belongs to.
accountName (SMB only)	The Active Directory user name of the client.
bucket	The S3 bucket name.
clientGid (NFS only)	The identifier of the group of the user accessing the object.
clientId (NFS only)	The identifier of the user accessing the object.
ctime	The time that the object's content or metadata was modified, set by the client.

Attribute	Definition
groupId	The identifier for group owner of the object.
fileSizeInBytes	The size of the file in bytes, set by the client at file creation time.
gateway	The Storage Gateway ID.
mtime	This time that the object's content was modified, set by the client.
newObjectName	The full path to the new object after it has been renamed.
objectName	The full path to the object.
objectType	Defines whether the object is a file or folder.
operation	The name of the object access operation.
ownerId	The identifier for the owner of the object.
securityDescriptor (SMB only)	Shows the discretionary access control list (DACL) set on an object, in SDDL format.
shareName	The name of the share that is being accessed.
source	The ID of the file share being audited.
sourceAddress	The IP address of file share client machine.
status	The status of the operation. Only success is logged (failures are logged with the exception of failures arising from permissions denied).
timestamp	The time that the operation occurred based on the OS timestamp of the gateway.
version	The version of the audit log format.

Attributes logged per operation

The following table describes the S3 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 - SMB only)	Write attributes (chown)	Write attributes (chmod)	Write attributes (chgrp)
access			X	X					X	
account main (SMB only)	X	X	X	X	X	X	X	X	X	X
account me (SMB only)	X	X	X	X	X	X	X	X	X	X
bucket	X	X	X	X	X	X	X	X	X	X
client (NFS only)	X	X	X	X	X	X		X	X	X
client (NFS only)	X	X	X	X	X	X		X	X	X

	Read data	Write data	Create folder	Create file	Rename file/ fold er	Delete file/ fold er	Write attribute s (change ACL - SMB only)	Write attribute s (chown)	Write attribute s (chmod)	Write attribute s (chgrp)
ctime			X	X						
groupID			X	X						
fileSize inBytes				X						
gateway	X	X	X	X	X	X	X	X	X	X
mtime			X	X						
newObjectName					X					
object e	X	X	X	X	X	X	X	X	X	X
object e	X	X	X	X	X	X	X	X	X	X
operat	X	X	X	X	X	X	X	X	X	X
ownerID			X	X				X		
securi escrip (SMB only)							X	X		

	Read data	Write data	Create folder	Create file	Rename file/ fold er	Delete file/ fold er	Write attribute s (change ACL - SMB only)	Write attribute s (chown)	Write attribute s (chmod)	Write attribute s (chgrp)
shareName	X	X	X	X	X	X	X	X	X	X
sourcePath	X	X	X	X	X	X	X	X	X	X
sourcePathPrefix	X	X	X	X	X	X	X	X	X	X
status	X	X	X	X	X	X	X	X	X	X
timestamp	X	X	X	X	X	X	X	X	X	X
versionId	X	X	X	X	X	X	X	X	X	X

Create a cache report for your S3 File Gateway

S3 File Gateway can generate a report of the metadata for files that are currently in the local upload cache for a specific file share. You can apply filters and additional criteria that determine which specific types of cached files appear in the report. You can use this report to help identify and resolve gateway issues. For example, if you have files failing upload from your gateway to Amazon S3, you can generate a report that lists the specific files that are failing to upload, and the reasons for upload failure. The report is a CSV file containing a list of files which match the set of filter parameters you specify. The output file is stored as an Amazon S3 object in a bucket location that you specify when you configure the report. To create a cache report using the AWS Storage Gateway API, see [StartCacheReport](#) in the *Storage Gateway API Reference*. To create a cache report in the Storage Gateway console, use the following procedure.

Prerequisites

- Your gateway must have `s3:PutObject` and `s3:AbortMultipartUpload` permissions for the Amazon S3 bucket where you want to store the cache report.
- No other cache reports can currently be in-progress for the file share.
- There must be fewer than 10 existing cache reports for the file share.
- The gateway must be online and connected to AWS.
- The gateway root disk must have at least 20GB of free space.

To create a cache report using the Storage Gateway console

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 **File shares**, and then choose the file share for which you want to create a cache report.
3. From the **Actions** drop-down menu, choose **Create cache report**.
4. For **Amazon S3 location**, enter the Amazon S3 bucket and prefix of the location where you want to save the completed cache report CSV file object in Amazon S3. To select the bucket and prefix from your existing Amazon S3 storage, choose **Browse S3**.
5. For **IAM role**, do one of the following to specify an IAM role that grants your File Gateway permissions to generate and store your cache report:
 - To specify an existing IAM role, choose a role from the drop-down list.
 - To manually create a new IAM role, choose **Create a role**, then create the new role using the IAM console.

Note

The IAM role you specify must have the following permissions to write objects to the report bucket **Amazon S3 location**, and to stop multipart uploads to the report bucket:

- `s3:PutObject`
- `s3:AbortMultipartUpload`

The role must also allow the `storagegateway.amazonaws.com` service to assume the role using the `sts:AssumeRole` action.

6. For **Report filter**, do one of the following to determine which files will be included in the cache report:
- To include all cached files that are currently failing upload to Amazon S3, choose **All files failing to upload**.
 - To include only those files that fail upload to Amazon S3 for a specific reason, choose **Specific upload failure reasons only**. Then, for **Failure reasons**, select one or more of the following reasons:
 - Inaccessible storage class — The gateway can't access the Amazon S3 storage class where the object is stored. For more information, see [Error: InaccessibleStorageClass](#).
 - Invalid object state — The state of the file on the gateway doesn't match its state in Amazon S3. For more information, see [Error: InvalidObjectState](#).
 - Object missing — The object has been deleted or moved in Amazon S3. For more information, see [Error: ObjectMissing](#).
 - S3 Access Denied — The Amazon S3 bucket access IAM role doesn't allow the gateway to perform the upload operation. For more information, see [Error: S3AccessDenied](#).

 **Note**

The **Files Failing Upload** flag is reset every 24 hours and during gateway reboot. If this report captures the files after the reset, but before they become flagged again, they will not be reported as **Files Failing Upload**.

7. For **Use a VPC endpoint to connect to S3?**, do one of the following to specify how your gateway will connect to the Amazon S3 bucket:
- To connect directly without using Amazon VPC, choose **Connect directly to the bucket**.
 - To browse a list of existing Amazon VPC endpoints, choose **Choose a VPC endpoint**, and then specify an endpoint from the **VPC endpoints** drop-down list that appears.
 - To specify an existing Amazon VPC endpoint by its DNS name, choose **Input a VPC endpoint DNS name**, and then enter the DNS name in the **VPC endpoint DNS name** field that appears.

Note

If your file share uses a VPC endpoint to connect to Amazon S3 for normal operations, we recommend using the same VPC when you configure your cache report. Cache report creation will fail if the gateway can't connect to the Amazon S3 bucket for any reason, including invalid VPC configuration.

8. (Optional) Under **Tags - optional**, choose **Add new tag**, then enter a **Key** and **Value** for your cache report.

A tag is a case-sensitive key-value pair that helps you categorize your Storage Gateway resources. Adding tags can make filtering and searching for your cache report easier. You can repeat this step to add up to 50 tags.

9. Choose **Create report** when finished.

Storage Gateway begins generating the report. You can check progress and view status on the **Cache reports** tab of the details page for the file share.

View and manage cache reports for your S3 File Gateway

Cache reports list files that are currently in the local cache for a specific file share, according to filters and criteria that you specify. You can view a list of existing cache reports for a specific file share, check report progress and status, and delete reports you no longer need using the AWS Storage Gateway API or the Storage Gateway console.

To manage cache reports using the API, see the following sections in the *Storage Gateway API Reference*:

- [ListCacheReports](#)
- [DescribeCacheReport](#)
- [CancelCacheReport](#)
- [DeleteCacheReport](#)

To manage cache reports in the Storage Gateway console, use the following procedure.

To manage cache reports using the Storage Gateway console

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 **File shares**, and then choose the file share for which you want to manage cache reports.
3. On the **Details** page for the file share, choose the **Cache reports** tab. This tab lists the existing cache reports for the file share, and provides information about status, progress, and the object path where the report file is stored in Amazon S3.
4. Do one of the following:
 - To view additional details for a specific report, such as the report ARN and associated tags, choose a report from the **Report ID** column.
 - To specify multiple reports to manage simultaneously, select the reports using the checkbox column.
5. To manage one or more reports, choose one of the following from the **Actions** drop-down menu:
 - Delete cache report — This deletes the record of the cache report from the Storage Gateway database. Delete records for obsolete cache reports to make room for new reports. Each file share can have up to 10 existing cache reports at any time.

Note

Deleting the cache report record using this procedure **does not** delete the report file object from Amazon S3.

- Cancel report — This cancels a report that is currently in-progress. Cancel an in-progress report if you made a mistake during report configuration, or if the report takes an unusually long time to complete. Confirm the cancellation when prompted.

Note

Completion times can vary significantly depending on the number of files in the cache. Typically, most reports complete within 5 minutes.

The Storage Gateway console displays a message indicating the result of the cancellation or delete action.

Understanding the information provided in S3 File Gateway cache reports

Cache reports list files that are currently in the local cache for a specific file share, according to filters and criteria that you specify. Each cache report includes the following information:

- **Bucket** — The Amazon S3 bucket or Access Point that is associated with the file share.
- **S3ObjectKey** — The Amazon S3 object that stores the data and metadata for this file. This object has the latest data that has been uploaded to S3, but it could be missing data which is failing to upload to S3.
- **FilePath** — The file path for the file entry in the gateway cache. This is where you can find the file when mounting and browsing the file share.
- **RenamedTo** — The new path of a renamed file. When you rename a file on your file share, the gateway needs to track both the old and new locations of the file. This field shows where the file was moved to, helping you track file rename operations - even if a file has been renamed multiple times. This information is particularly useful when you need to understand how files in your file share correspond to objects in your Amazon S3 bucket.

The following example shows the cache report entries for a complex scenario involving a file being overwritten directly in Amazon S3, while also being renamed through File Gateway. In this scenario, the gateway uploads file `A.txt` to S3, and then evicts the file contents to make space in the local cache. The associated S3 object is then overwritten directly in S3—not through an action taken by the gateway—which results in an `InvalidObjectState` due to the mismatch between the S3 object and what the gateway expects. At the same time, file `A.txt` was renamed to `B.txt` through the gateway.

Bucket	S3ObjectKey	FilePath	Renamed	Type	IsDirty	IsDataDirty	IsDeleted	IsFailingToUpload	UploadError	SizeInBytes	IsWholeFileInCache
sample-ket- iad	A.txt	/B.txt		FILE	TRUE	FALSE	FALSE	TRUE	Invalid objectSta	4	FALSE
sample-ket- iad	A.txt	/A.txt	/B.txt	FILE	TRUE	FALSE	TRUE	FALSE		4	FALSE

- **Type** — Denotes whether the entry is for a FILE or DIRECTORY.
- **IsDirty** — Reports TRUE if there is any type of change to the file which have not been uploaded to Amazon S3. This includes changes to metadata such as file name and read/write permissions, even if the file's data has not changed.
- **IsDataDirty** — Reports TRUE if there are changes to the file's data which have not been uploaded to Amazon S3.
- **IsDeleted** — Reports TRUE if the file was deleted on the gateway. If a file is marked as deleted, then it will always be marked as dirty.
- **IsFailingToUpload** — Reports TRUE if there is a problem uploading the file to Amazon S3. This status resets every 24 hours to allow the gateway to retry the upload and check whether the issue has been resolved. The gateway rejects any new write operations for a file that is failing to upload. If the gateway does not have the entire file in cache, then it also rejects read operations.
- **UploadError** — The error that is preventing the file from uploading to Amazon S3. For more information and recommended steps to resolve these errors, see [Troubleshooting: File Gateway issues](#).
- **SizeInBytes** — The total size of the file.
- **IsWholeFileInCache** — Reports TRUE if all of the file's data is currently stored in the gateway cache. If this is TRUE for a file failing to upload to Amazon S3, then the gateway will allow the file to be read.

Maintaining your gateway

Maintaining your Amazon S3 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 S3 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 S3 File Gateway with a new instance](#) - Learn how to replace your S3 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 [Release Notes for Gateway Appliance Software](#).

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 to 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
<p>Get information about your network adapter</p>	<p>Enter the corresponding numeral to select Describe Adapter.</p> <p>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:</p> <ul style="list-style-type: none">• Media access control (MAC) address• IP address• Netmask• Gateway IP address• DHCP enabled status <p>You use the adapter names listed here when you configure a static IP address or when you set your gateway's default adapter.</p>
<p>Configure DHCP routing</p>	<p>Enter the corresponding numeral to select Configure DHCP.</p> <p>You are prompted to configure the network interface to use DHCP.</p>

To Perform This Task	Do This
<p>Configure a static IP address for your gateway</p>	<p>Enter the corresponding numeral to select Configure Static IP.</p> <p>You are prompted to enter the following information to configure a static IP:</p> <ul style="list-style-type: none">• Network adapter name• IP address• Netmask• Default gateway address• Primary Domain Name Service (DNS) address• Secondary DNS address <div data-bbox="829 1209 1511 1621"><p> Important</p><p>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.</p></div> <p>If your gateway uses more than one network interface, you must set all active interfaces to use DHCP or static IP addresses.</p>

To Perform This Task	Do This
	<p>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.</p> <p>If both interfaces are initially set to use static IP addresses and you then set the gateway to use DHCP, both interfaces use DHCP.</p>
Configure a hostname for your gateway	<p>Enter the corresponding numeral to select Configure Hostname.</p> <p>You are prompted to choose whether the gateway will use a static hostname that you specify, or acquire one automatically through DHCP or rDNS.</p> <p>If you select Static, you are prompted to provide a static hostname, such as <code>testgateway.example.com</code>. Enter <code>y</code> to apply the configuration.</p> <div data-bbox="829 1276 1510 1734"><p> Note</p><p>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.</p></div>

To Perform This Task	Do This
View your gateway's hostname configuration	<p>Enter the corresponding numeral to select View Hostname Configuration.</p> <p>Your gateway's hostname, acquisition mode, domain, and Active Directory realm are displayed.</p>
Reset all your gateway's network configuration to DHCP	<p>Enter the corresponding numeral to select Reset all to DHCP.</p> <p>All network interfaces are set to use DHCP.</p> <div data-bbox="829 800 1511 1209"><p> Important</p><p>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.</p></div>
Set your gateway's default route adapter	<p>Enter the corresponding numeral to select Set Default Adapter.</p> <p>The available adapters for your gateway are shown, and you are prompted to choose one of the adapters—for example, eth0.</p>

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

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.	<p>Enter the corresponding numeral to select View and Synchronize System Time.</p> <p>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.</p> <p>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</p>

To Perform This Task	Do This
	<p>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.</p> <p>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 information, see Synchronize VM time with VMware host time.</p> <p>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.</p> <p>For a gateway deployed on KVM, you can check and synchronize the VM time using <code>virsh</code> command line interface for KVM.</p>
Edit your NTP server configuration	<p>Enter the corresponding numeral to select Edit NTP Configuration.</p> <p>You are prompted to provide a preferred and a secondary NTP server.</p>
View your NTP server configuration	<p>Enter the corresponding numeral to select View NTP Configuration.</p> <p>Your NTP server configuration is displayed.</p>

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 troubleshooting.
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
	<div> local console menu option. For instructions, see Configuring your gateway network settings. </div>
ip	<p>Show / manipulate routing, devices, and tunnels.</p> <div> 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. </div>
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
tcptracert	Collect traceroute output on TCP traffic to a destination.
sslcheck	Returns output with certificate issuer <div>  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. </div>

- 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 recommended 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 troubleshooting.
exit	Return to Configuration menu.
h	Display available command list.
ifconfig	<div>View or configure network interfaces.</div> <div> 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.</div>
ip	<div>Show / manipulate routing, devices, and tunnels.</div> <div> Note We recommend configuring network or IP settings using the Storage Gateway console or the dedicated</div>

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.

- 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

- Log in to the local console on your Amazon EC2 File Gateway. For instructions, see [Logging in to your Amazon EC2 gateway local console](#).
- 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 acquire a hostname automatically through DHCP 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, acquisition 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 S3 File Gateway with a new instance

You can replace an existing S3 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.

There are two methods to replace an existing S3 File Gateway. The following table describes the benefits and drawbacks of each method. Using this information, select the method best suited for your gateway environment, then refer to the procedure steps in the corresponding section below.

Note

If you need to log into your new Storage Gateway's local console to complete either method, the default username is *admin*, and the default password is *password*.

	Method 1: Migrate cache disk and Gateway ID to replacement instance	Method 2: Replacement instance with empty cache disk and new Gateway ID
Cache disk data	Data on the cache disk is preserved. This method is useful if your gateway has a large cache disk, or if your applications are sensitive to the delay caused by out-of-cache read operations.	Data in cache is downloaded from the AWS cloud. This method is optimal for write-heavy workloads, if your applications can tolerate the delay caused by out-of-cache reads.
Down time	Your gateway will be offline for 1-2 hours during the migration process.	File shares are always available, but clients will experience short cutover downtime when switching from one file share to another during the transition to the new instance.

	Method 1: Migrate cache disk and Gateway ID to replacement instance	Method 2: Replacement instance with empty cache disk and new Gateway ID
		<p>Note</p> <p>Writing to one Amazon S3 bucket from two file shares simultaneously is <i>not supported</i>, so all clients must be remapped from one share to the other simultaneously, rather than gradually.</p>
Gateway ID	The new gateway inherits the Gateway ID from the gateway it replaces.	The existing gateway and replacement gateway have separate, unique Gateway IDs.

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.

Method 1: Migrate cache disk and Gateway ID to replacement instance

To migrate your S3 File Gateway's cache disk and Gateway ID to a replacement instance:

1. Stop any applications that are writing to the existing S3 File Gateway.
2. Verify that the CachePercentDirty metric on the **Monitoring** tab for the existing S3 File Gateway is 0.
3. Shut down the existing S3 File Gateway by powering off the host virtual machine (VM) using its hypervisor controls.

For more information about shutting down an Amazon EC2 instance, see [Stop and start your instance](#) in the *Amazon EC2 User Guide*.

For more information about shutting down a KVM, VMware, or Hyper-V VM, see your hypervisor documentation.

4. Detach all disks, including the root disk and cache disks from the old gateway VM.

 **Note**

Make a note of the root disk's volume ID, as well as the gateway ID associated with that root disk. You will need to detach this disk from the new Storage Gateway hypervisor in a later step.

If you are using an Amazon EC2 instance as the VM for your S3 File Gateway, see [Detach an Amazon EBS volume from a Windows instance](#) or [Detach an Amazon EBS volume from a Linux instance](#) in the *Amazon EC2 User Guide*.

For information about detaching disks from a KVM, VMware, or Hyper-V VM, see the documentation for your hypervisor.

5. Create a new AWS Storage Gateway hypervisor VM instance, but don't activate it as a gateway. In a later step, this new VM will assume the identity of the old gateway.

For more information about creating a new Storage Gateway hypervisor VM, see [Choosing a Host Platform and Downloading the VM](#).

 **Note**

Do not add cache disks for the new VM. This VM will use the same cache disks that were used by the old VM.

6. Configure your new Storage Gateway VM to use the same network settings as the old VM.

The default network configuration for the gateway is Dynamic Host Configuration Protocol (DHCP). With DHCP, your gateway is automatically assigned an IP address.

If you need to manually configure a static IP address for your gateway VM, see [Configuring network parameters](#).

If your gateway VM must use a Socket Secure version 5 (SOCKS5) proxy to connect to the internet, see [Routing your gateway deployed on EC2 through an HTTP proxy](#).

7. Start the new Storage Gateway VM.
8. Attach the disks that you detached from the old gateway VM to the new gateway VM. Do not detach the existing root disk from the new gateway VM.

 **Note**

To migrate successfully, all disks must remain unchanged. Changing the disk size or other values causes inconsistencies in metadata that prevent successful migration.

9. Initiate the gateway migration process by connecting to the new VM with a URL that uses the following format:

`http://your-VM-IP-address/migrate?gatewayId=your-gateway-ID`

You can use the same IP address for the new gateway VM that you used for the old gateway VM. To use the same IP, delete the old gateway first and then proceed from step 5. Your URL should look similar to the following example:

`http://198.51.100.123/migrate?gatewayId=sgw-12345678`

Use this URL from a browser, or from the command line using cURL.

When the gateway migration initiates successfully, the following message appears:

Successfully imported Storage Gateway information. Please refer to Storage Gateway documentation to perform the next steps to complete the migration.

10. Wait for the gateway status to show as **Running** in the AWS Storage Gateway console. Depending on available bandwidth, this can take up to 10 minutes.
11. Stop the new Storage Gateway VM.
12. Detach the old gateway's root disk, whose volume ID you noted previously, from the new gateway.
13. Start the new Storage Gateway VM.
14. If your gateway was joined to an Active Directory domain, re-join the domain. For instructions, see [Using Active Directory to authenticate users](#).

Note

You must complete this step even if the status of the S3 File Gateway appears as **Joined**.

15. Confirm that your shares are available at the new gateway VM's IP address, then delete the old gateway VM.

Warning

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

For more information about deleting an Amazon EC2 instance, see [Terminate your instance](#) in the *Amazon EC2 User Guide*. For more information about deleting a KVM, VMware, or Hyper-V VM, see the documentation for your hypervisor.

Method 2: Replacement instance with empty cache disk and new Gateway ID

To set up a replacement S3 File Gateway instance with empty cache disk and new Gateway ID:

1. Stop any applications that are writing to the existing S3 File Gateway. Verify that the CachePercentDirty metric on the **Monitoring** tab is 0 before you set up file shares on the new gateway.
2. Use the AWS Command Line Interface (AWS CLI) to gather and save the configuration information about your existing S3 File Gateway and file shares by doing the following:
 - a. Save the gateway configuration information for the S3 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 S3 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 metadata about the SMB file share, such as its domain name, Microsoft Active Directory status, whether the guest password is set, and the type of security strategy.

- c. Save file share information for each SMB and Network File System (NFS) file share of the S3 File Gateway:
 - Use the following command for SMB file shares.

```
aws storagegateway describe-smb-file-shares --file-share-arn-list
"arn:aws:storagegateway:us-east-2:123456789012:share/share-987A654B"
```

This command outputs a JSON block that contains metadata about the SMB file share, such as its name, storage class, status, IAM role Amazon Resource Name (ARN), a list of clients that are allowed to access the S3 File Gateway, and the path used by the SMB client to identify the mount point.

- Use the following command for NFS file shares.

```
aws storagegateway describe-nfs-file-shares --file-share-arn-list
"arn:aws:storagegateway:us-east-2:123456789012:share/share-321A978B"
```

This command outputs a JSON block that contains metadata about the NFS file share, such as its name, storage class, status, IAM role ARN, a list of clients that are allowed to access the S3 File Gateway, and the path used by the NFS client to identify the mount point.

3. Create a new S3 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 shares for the new gateway with the same settings and configuration as the file shares 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 shares to the new file shares 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 S3 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 share 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. If you don't need the data in Amazon S3 buckets associated with the file shares, you can choose to remove your Amazon S3 buckets. For instructions, see [Deleting your bucket](#).

Performance and optimization

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

Topics

- [Basic performance guidance for S3 File Gateway](#)
- [Performance guidance for gateways with multiple file shares](#)
- [Maximizing S3 File Gateway throughput](#)
- [Optimizing S3 File Gateway for SQL Server database backups](#)

Basic performance guidance for S3 File Gateway

In this section, you can find guidance for provisioning hardware for your S3 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](#).

For Amazon EC2 instances, if you have more than 5 million objects in your S3 bucket and you are using a General Purposes SSD volume, a minimum root EBS volume of 350 GiB is needed for acceptable performance of your gateway during start up. For information about how to increase your volume size, see [Modifying an EBS volume using elastic volumes \(console\)](#).

The suggested size limit for individual directories in the file shares 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 shares that are served from cache. *Cache miss* read operations are reads from the file shares that are served from Amazon S3.

The following tables show example S3 File Gateway configurations.

S3 File Gateway performance on Linux clients

Example Configurations	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 disk: 512 GiB cache, io1, 1,500 provisioned IOPS Minimum network performance: 10 Gbps CPU: 16 vCPU RAM: 32 GB NFS protocol recommended for Linux	NFSv3 - 1 thread	110 MiB/sec (0.92 Gbps)	590 MiB/sec (4.9 Gbps)	310 MiB/sec (2.6 Gbps)
	NFSv3 - 8 threads	160 MiB/sec (1.3 Gbps)	590 MiB/sec (4.9 Gbps)	335 MiB/sec (2.8 Gbps)
	NFSv4 - 1 thread	130 MiB/sec (1.1 Gbps)	590 MiB/sec (4.9 Gbps)	295 MiB/sec (2.5 Gbps)
	NFSv4 - 8 threads	160 MiB/sec (1.3 Gbps)	590 MiB/sec (4.9 Gbps)	335 MiB/sec (2.8 Gbps)
Storage Gateway Hardware Appliance	SMBV3 - 1 thread	115 MiB/sec (1.0 Gbps)	325 MiB/sec (2.7 Gbps)	255 MiB/sec (2.1 Gbps)
	SMBV3 - 8 threads	190 MiB/sec (1.6 Gbps)	590 MiB/sec (4.9 Gbps)	335 MiB/sec (2.8 Gbps)
Minimum network performance: 10 Gbps	NFSv3 - 1 thread	265 MiB/sec (2.2 Gbps)	590 MiB/sec (4.9 Gbps)	310 MiB/sec (2.6 Gbps)
	NFSv3 - 8 threads	385 MiB/sec (3.1 Gbps)	590 MiB/sec (4.9 Gbps)	335 MiB/sec (2.8 Gbps)
	NFSv4 - 1 thread	310 MiB/sec (2.6 Gbps)	590 MiB/sec (4.9 Gbps)	295 MiB/sec (2.5 Gbps)

Example Configurations	Protocol	Write throughput (file sizes 1 GB)	Cache hit read throughput	Cache miss read throughput
	NFSv4 - 8 threads	385 MiB/sec (3.1 Gbps)	590 MiB/sec (4.9 Gbps)	335 MiB/sec (2.8 Gbps)
	SMBV3 - 1 thread	275 MiB/sec (2.4 Gbps)	325 MiB/sec (2.7 Gbps)	255 MiB/sec (2.1 Gbps)
	SMBV3 - 8 threads	455 MiB/sec (3.8 Gbps)	590 MiB/sec (4.9 Gbps)	335 MiB/sec (2.8 Gbps)
Root disk: 80 GB, io1 SSD, 4,000 IOPS Cache disk: 4 x 2 TB NVME cache disks	NFSv3 - 1 thread	300 MiB/sec (2.5 Gbps)	590 MiB/sec (4.9 Gbps)	325 MiB/sec (2.7 Gbps)
	NFSv3 - 8 threads	585 MiB/sec (4.9 Gbps)	590 MiB/sec (4.9 Gbps)	580 MiB/sec (4.8 Gbps)
	NFSv4 - 1 thread	355 MiB/sec (3.0 Gbps)	590 MiB/sec (4.9 Gbps)	340 MiB/sec (2.9 Gbps)
Minimum network performance: 10 Gbps CPU: 32 vCPU RAM: 244 GB NFS protocol recommended for Linux	NFSv4 - 8 threads	575 MiB/sec (4.8 Gbps)	590 MiB/sec (4.9 Gbps)	575 MiB/sec (4.8 Gbps)
	SMBV3 - 1 thread	230 MiB/sec (1.9 Gbps)	325 MiB/sec (2.7 Gbps)	245 MiB/sec (2.0 Gbps)
	SMBV3 - 8 threads	585 MiB/sec (4.9 Gbps)	590 MiB/sec (4.9 Gbps)	580 MiB/sec (4.8 Gbps)

File Gateway performance on Windows clients

Example Configurations	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 disk: 512 GiB cache, io1, 1,500 provisioned IOPS Minimum network performance: 10 Gbps CPU: 16 vCPU RAM: 32 GB SMB protocol recommended for Windows	SMBV3 - 1 thread	150 MiB/sec (1.3 Gbps)	180 MiB/sec (1.5 Gbps)	20 MiB/sec (0.2 Gbps)
	SMBV3 - 8 threads	190 MiB/sec (1.6 Gbps)	335 MiB/sec (2.8 Gbps)	195 MiB/sec (1.6 Gbps)
	NFSv3 - 1 thread	95 MiB/sec (0.8 Gbps)	130 MiB/sec (1.1 Gbps)	20 MiB/sec (0.2 Gbps)
	NFSv3 - 8 threads	190 MiB/sec (1.6 Gbps)	330 MiB/sec (2.8 Gbps)	190 MiB/sec (1.6 Gbps)
Storage Gateway Hardware Appliance Minimum network performance: 10 Gbps	SMBV3 - 1 thread	230 MiB/sec (1.9 Gbps)	255 MiB/sec (2.1 Gbps)	20 MiB/sec (0.2 Gbps)
	SMBV3 - 8 threads	835 MiB/sec (7.0 Gbps)	475 MiB/sec (4.0 Gbps)	195 MiB/sec (1.6 Gbps)
	NFSv3 - 1 thread	135 MiB/sec (1.1 Gbps)	185 MiB/sec (1.6 Gbps)	20 MiB/sec (0.2 Gbps)
	NFSv3 - 8 threads	545 MiB/sec (4.6 Gbps)	470 MiB/sec (4.0 Gbps)	190 MiB/sec (1.6 Gbps)

Example Configurations	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 disk: 4 x 2 TB NVME cache disks Minimum network performance: 10 Gbps CPU: 32 vCPU RAM: 244 GB SMB protocol recommended for Windows	SMBV3 - 1 thread	230 MiB/sec (1.9 Gbps)	265 MiB/sec (2.2 Gbps)	30 MiB/sec (0.3 Gbps)
	SMBV3 - 8 threads	835 MiB/sec (7.0 Gbps)	780 MiB/sec (6.5 Gbps)	250 MiB/sec (2.1 Gbps)
	NFSv3 - 1 thread	135 MiB/sec (1.1 Gbps)	220 MiB/sec (1.8 Gbps)	30 MiB/sec (0.3 Gbps)
	NFSv3 - 8 threads	545 MiB/sec (4.6 Gbps)	570 MiB/sec (4.8 Gbps)	240 MiB/sec (2.0 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.

Performance guidance for gateways with multiple file shares

Amazon S3 File Gateway supports attaching up to 50 file shares to a single Storage Gateway appliance. By adding multiple file shares per gateway, you can support more users and workloads while managing fewer gateways and virtual hardware resources. In addition to other factors, the number of file shares managed by a gateway can affect that gateway's performance. This section

describes how gateway performance is expected to change depending on the number of attached file shares and recommends virtual hardware configurations to optimize performance for gateways that manage multiple shares.

In general, increasing the number of file shares managed by a single Storage Gateway can have the following consequences:

- Increased time required to restart the gateway.
- Increased utilization of virtual hardware resources such as vCPU and RAM.
- Decreased performance for data and metadata operations if virtual hardware resources become saturated.

The following table lists recommended virtual hardware configurations for gateways that manage multiple file shares:

File Shares Per Gateway	Recommended Gateway Capacity Setting	Recommended vCPU Cores	Recommended RAM	Recommended Disk Size			
1-10	Small	4 (EC2 instance type m4.xlarge or greater)	16 GiB	80 GiB			
10-20	Medium	8 (EC2 instance type m4.2xlarge or greater)	32 GiB	160 GiB			
20+	Large	16 (EC2 instance	64 GiB	240 GiB			

File Shares Per Gateway	Recommended Gateway Capacity Setting	Recommended vCPU Cores	Recommended RAM	Recommended Disk Size			
		type m4.4xlarge e or greater)					

In addition to the virtual hardware configurations recommended above, we recommend the following best practices for configuring and maintaining Storage Gateway appliances that manage multiple file shares:

- Consider that the relationship between the number of file shares and the demand placed on the gateway's virtual hardware is not necessarily linear. Some file shares might generate more throughput, and therefore more hardware demand than others. The recommendations in the preceding table are based on maximum hardware capacities and various file share throughput levels.
- If you find that adding multiple file shares to a single gateway reduces performance, consider moving the most active file shares to other gateways. In particular, if a file share is used for a very-high-throughput application, consider creating a separate gateway for that file share.
- We do not recommend configuring one gateway for multiple high-throughput applications and another for multiple low-throughput applications. Instead, try to spread high and low throughput file shares evenly across gateways to balance hardware saturation. To measure your file share throughput, use the `ReadBytes` and `WriteBytes` metrics. For more information, see [Understanding file share metrics](#).

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 significantly 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 recommend 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

Storage Gateway uses SSL/TLS (Secure Socket Layers/Transport Layer Security) to encrypt data that is transferred between your gateway appliance and AWS storage. By default, Storage Gateway uses Amazon S3-Managed encryption keys (SSE-S3) to server-side encrypt all data it stores in Amazon S3. You have an option to use the Storage Gateway API to configure your gateway to encrypt data stored in the cloud using server-side encryption with AWS Key Management Service (SSE-KMS) keys.

Encrypting a file share

You can configure the file shares on your S3 File Gateway to encrypt stored objects with AWS KMS-managed keys by using SSE-KMS or DSSE-KMS. For information about supported file share encryption methods, see [Encrypt objects stored by File Gateway in Amazon S3](#).

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 S3.

- 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](#)
- [Authenticating with identities](#)
- [Managing access using policies](#)
- [How AWS Storage Gateway works with IAM](#)
- [Identity-based policy examples for AWS Storage Gateway](#)
- [Troubleshooting AWS Storage Gateway identity and access](#)
- [Using tags to control access to your gateway and resources](#)
- [Using Windows ACLs to limit SMB file share access](#)

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"
        }
      }
    }
  ]
}
```

```
}  
  }  
    ]  
      }
```

Using Windows ACLs to limit SMB file share access

Amazon S3 file gateway supports two different methods for controlling access to files and directories that are stored through an SMB file share: POSIX permissions, or Windows ACLs.

This section describes how to use Microsoft Windows access control lists (ACLs) on SMB file shares that use Microsoft Active Directory (AD) for authentication. By using Windows ACLs, you can set fine-grained permissions on files and folders in your SMB file share.

Following are some important characteristics of Windows ACLs on SMB file shares:

- Windows ACLs are selected by default for SMB file shares when your file gateway is joined to an Active Directory domain.
- When ACLs are activated, the ACL information is persisted in Amazon S3 object metadata.
- The gateway preserves up to 10 ACLs per file or folder.
- When you use an SMB file share with ACLs activated to access S3 objects created outside your gateway, the objects inherit ACLs' information from the parent folder.

Note

The default root ACL for an SMB file share gives full access to everyone, but you can change the permissions of the root ACL. You can use root ACLs to control access to the file share. You can set who can mount the file share (map the drive) and what permissions the user gets to the files and folders recursively in the file share. However, we recommend that you set this permission on the top-level folder in the S3 bucket so that your ACL is persisted.

You can turn on Windows ACLs when you create a new SMB file share by using the [CreateSMBFileShare](#) API operation. Or you can turn on Windows ACLs on an existing SMB file share by using the [UpdateSMBFileShare](#) API operation.

Activating Windows ACLs on a New SMB File Share

Take the following steps to activate Windows ACLs on a new SMB file share.

To activate Windows ACLs when creating a new SMB file share

1. Create a file gateway if you don't already have one. For more information, see [Creating your gateway](#).
2. If the gateway is not joined to a domain, add it to a domain. For more information, see [Using Active Directory to authenticate users](#).
3. Create an SMB file share. For more information, see
4. Activate Windows ACLs on the file share from the Storage Gateway console.

To use the Storage Gateway Console, do the following:

- a. Choose the file share and choose **Edit file share**.
 - b. For the **File/directory access controlled by** option, choose **Windows Access Control List**.
5. (Optional) Add an admin user to the [AdminUsersList](#), if you want the admin user to have privileges to update ACLs on all files and folders in the file share.

Note

If you have configured the **Allowed and Denied Users and Groups** lists in the SMB file share's settings, then ACLs will not grant any access that overrides those lists.

The **Allowed and Denied Users and Groups** lists are evaluated before ACLs, and control which users can mount or access the file share. If any users or groups are placed on the **Allowed** list, the list is considered active, and only those users can mount the file share.

After a user has mounted a file share, ACLs then provide more granular protection that controls which specific files or folders the user can access.

6. Update the ACLs for the parent folders under the root folder. To do this, use Windows File Explorer to configure the ACLs on the folders in the SMB file share.

 **Note**

If you configure the ACLs on the root instead of the parent folder under root, the ACL permissions aren't persisted in Amazon S3.

We recommend setting ACLs at the top-level folder under the root of your file share, instead of setting ACLs directly at the root of the file share. This approach persists the information as object metadata in Amazon S3.

7. Turn on inheritance as appropriate.

 **Note**

You can turn on inheritance for file shares created after May 8, 2019.

If you turn on inheritance and update the permissions recursively, Storage Gateway updates all the objects in the S3 bucket. Depending on the number of objects in the bucket, the update can take a while to complete.

Activating Windows ACLs on an Existing SMB File Share

Take the following steps to activate Windows ACLs on an existing SMB file share that has POSIX permissions.

To activate Windows ACLs on an existing SMB file share using the Storage Gateway Console

1. Choose the file share and choose **Edit file share**.
2. For the **File/directory access controlled by** option, choose **Windows Access Control List**.
3. Turn on inheritance as appropriate.

 **Note**

We don't recommend setting the ACLs at the root level, because if you do this and delete your gateway, you need to reset the ACLs again.

If you turn on inheritance and update the permissions recursively, Storage Gateway updates all the objects in the S3 bucket. Depending on the number of objects in the bucket, the update can take a while to complete.

Limitations When Using Windows ACLs

Keep the following limitations in mind when using Windows ACLs to control access to SMB file shares:

- Windows ACLs are only supported on file shares that use Active Directory for authentication when you use Windows SMB clients to access the file shares.
- File gateways support a maximum of 10 ACL entries for each file and directory.
- File gateways don't support Audit and Alarm entries, which are system access-control list (SACL) entries. File gateways support Allow and Deny entries, which are discretionary access control list (DACL) entries.
- File gateways don't support Advanced Access Control Entry (ACE) permissions.
- The root ACL settings of SMB file shares are only on the gateway, and the settings are persisted across gateway updates and restarts.

Note

If you configure the ACLs on the root instead of the parent folder under the root, the ACL permissions aren't persisted in Amazon S3.

Given these conditions, make sure to do the following:

- If you configure multiple gateways to access the same Amazon S3 bucket, configure the root ACL on each of the gateways to keep the permissions consistent.
- If you delete a file share and recreate it on the same Amazon S3 bucket, make sure that you use the same set of root ACLs.

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": "AIDAI5AUPEBH2M7JTNC",
```

```

        "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": "AIDAI5AUPEPBH2M7JTNVC",
      "arn": "arn:aws:iam::111122223333:user/StorageGateway-
team/JohnDoe",

```

```

AKIAIOSFODNN7EXAMPLE",
    "accountId":" 111122223333", " accessKeyId ":"
    " 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 ":"
6U2N42CU37KA08BG6V1I23FRSJ1Q8GLLE1QEU3KPGG6F0KSTAUU0 ",
    " 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 shares, high availability, data recovery, and security. 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: file share issues](#) - Learn about actions you can take if you experience unexpected issues with your file share.
- [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, 0 = Amazon, CN = Amazon Root CA 1
verify return:1
depth=1 C = US, 0 = 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, 0 = Amazon, OU = Server CA 1B, CN = Amazon
1 s:C = US, 0 = Amazon, OU = Server CA 1B, CN = Amazon
  i:C = US, 0 = Amazon, CN = Amazon Root CA 1
2 s:C = US, 0 = 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, 0 = "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) endpoint 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.	<p>Use the hypervisor client to connect to your host to find the gateway IP address.</p> <ul style="list-style-type: none">• 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. <p>If you are still having trouble finding the gateway IP address:</p> <ul style="list-style-type: none">• 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.	<ul style="list-style-type: none">• 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 activation fails when you click the Proceed to Activation	<ul style="list-style-type: none">• Check that the gateway VM can be accessed by pinging the VM from your client.

Issue	Action to Take
<p>button in the Storage Gateway Management Console.</p>	<ul style="list-style-type: none"> • 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.
<p>You need to improve bandwidth between your gateway and AWS.</p>	<p>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 <code>CloudBytesDownloaded</code> and <code>CloudBytesUploaded</code> 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.</p>

Issue	Action to Take
Throughput to or from your gateway drops to zero.	<ul style="list-style-type: none"> 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. <p>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.</p>
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.

Troubleshooting: Security scans show open NFS ports

Certain NFS ports are enabled by default, even on gateways that you only use with SMB file shares. If you use third-party security software such as Qualys to scan the network where your File Gateway is deployed, the scan results might report these open NFS ports as a potential security vulnerability. If you only use your gateway with SMB file shares and you want to disable the unused NFS ports for security reasons, use the following procedure:

To disable NFS ports on a File Gateway:

1. Access the gateway local console command prompt using the procedure outlined in <https://docs.aws.amazon.com/filegateway/latest/files3/MaintenanceGatewayConsole-fgw.html>.
2. Enter the following commands to disable NFS traffic:

```
iptables -I INPUT -p udp -m udp --dport 111 -j DROP
iptables -I INPUT -p udp -m udp --dport 2049 -j DROP
iptables -I INPUT -p udp -m udp --dport 20048 -j DROP
iptables -I INPUT -p tcp -m tcp --dport 111 -j DROP
iptables -I INPUT -p tcp -m tcp --dport 2049 -j DROP
iptables -I INPUT -p tcp -m tcp --dport 20048 -j DROP
```

3. Enter the following command to confirm that the blocked NFS ports appear in the IP tables:

```
iptables -n -L -v --line-numbers
```

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-channel1** 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-channel1**. 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
<p>You try to import a gateway and receive the following error message:</p> <p>"A server error occurred while attempting to import the virtual machine. Import failed. Unable to find virtual machine import files under location [...]. You can import a virtual machine only if you used Hyper-V to create and export it."</p>	<p>This error can occur for the following reasons:</p> <ul style="list-style-type: none">• 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 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.

Issue	Action to Take
	<p>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.</p>
<p>You try to import a gateway and receive the following error message:</p> <p>"A server error occurred while attempting to import the virtual machine. Import failed. Import task failed to copy file from [...]: The file exists. (0x80070050)"</p>	<p>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.</p>
<p>You try to import a gateway and receive the following error message:</p> <p>"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."</p>	<p>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.</p>

Issue	Action to Take
<p>You try to start a gateway VM and receive the following error message:</p> <p>"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 [...])"</p>	<p>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.</p> <p>For more information about the requirements for Storage Gateway, see File Gateway setup requirements.</p>
<p>You try to start a gateway VM and receive the following error message:</p> <p>"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)"</p>	<p>This error is likely caused by a RAM discrepancy between the required RAM for the gateway and the available RAM on the host.</p> <p>For more information about the requirements for Storage Gateway, see File Gateway setup requirements.</p>

Issue	Action to Take
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 <code>hyperv-server</code> , then you can use the following UNC path <code>\\hyperv-server\c\$</code> , which assumes that the name <code>hyperv-server</code> 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.
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 S3 File Gateway](#).

For information about using ephemeral storage, see [Using ephemeral storage with EC2 gateways](#).

Topics

- [Your gateway activation hasn't occurred after a few moments](#)

- [You can't find your EC2 gateway instance in the instance list](#)
- [You want to connect to your gateway instance using the Amazon EC2 serial console](#)
- [You want Support to help troubleshoot your Amazon EC2 gateway](#)

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 determine the service IP address](#)
- [How do you perform a factory reset?](#)
- [How do you perform a remote restart?](#)
- [Where do you obtain Dell iDRAC support?](#)
- [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: 1344 \(0x00000540\)](#)
- [Error: GatewayClockOutOfSync](#)
- [Error: InaccessibleStorageClass](#)
- [Error: InvalidObjectState](#)
- [Error: ObjectMissing](#)
- [Error: RoleTrustRelationshipInvalid](#)
- [Error: S3AccessDenied](#)
- [Error: DroppedNotifications](#)
- [Notification: HardReboot](#)

- [Notification: Reboot](#)
- [Troubleshooting: Security scans show open NFS ports](#)
- [Troubleshooting: Using CloudWatch metrics](#)

Error: 1344 (0x00000540)

While migrating files to Amazon S3 you may encounter an ERROR 1344 (0x00000540) if you are trying to copy files with more than 10 Access Control Entries (ACEs) into Amazon S3. Access Control Entries are listed in the Access Control List (ACL).

The Amazon S3 File Gateway can only preserve 10 ACE entries per given file or folder.

To resolve an Error 1344: Copying NTFS Security to Destination Directory.

Reduce the number of entries in Windows Permissions for files or folders that contain more than 10 entries. A common approach is to create a group containing the full list of entries, then replacing the list of entries with that single group. Once the number of entries is less the 10, you can retry copying the files or folders to the gateway.

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: InaccessibleStorageClass

You can get an InaccessibleStorageClass error when an object has moved out of the Amazon S3 Standard storage class.

Your File Gateway usually encounters this error when it tries to either upload an object to or read an object from the Amazon S3 bucket. Generally, this error means the object has moved to Amazon S3 Glacier and is in either the S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive storage class.

Your S3 File Gateway can generate a cache report that lists all files in the gateway cache that are currently failing to upload to Amazon S3 due to this error. The information in this report can help you work with Support to resolve issues with your gateway, Amazon S3, or IAM configuration. For more information, see [Create a cache report](#).

To resolve an `InaccessibleStorageClass` error

- Restore the object from the S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive storage class back to its original storage class in S3.

If you restore the object to the S3 bucket to fix an upload error, the file is eventually uploaded. If you restore the object to fix a read error, the File Gateway's SMB or NFS client can then read the file.

Error: `InvalidObjectState`

You can get an `InvalidObjectState` error when a writer other than the specified File Gateway modifies the specified file in the specified Amazon S3 bucket. As a result, the state of the file for the File Gateway doesn't match its state in Amazon S3. Any subsequent uploads of the file to Amazon S3 or retrievals of the file from Amazon S3 fail.

Your S3 File Gateway can generate a cache report that lists all files in the gateway cache that are currently failing to upload to Amazon S3 due to this error. The information in this report can help you work with Support to resolve issues with your gateway, Amazon S3, or IAM configuration. For more information, see [Create a cache report](#).

To resolve an `InvalidObjectState` error

If the operation that modifies the file is `S3Upload` or `S3GetObject`, do the following:

1. Save the latest copy of the file to the local file system of your SMB or NFS client (you need this file copy in step 4). If the version of the file in Amazon S3 is the latest, download that version. You can do this using the AWS Management Console or AWS CLI.
2. Delete the file in Amazon S3 using the AWS Management Console or AWS CLI.
3. Delete the file from the File Gateway using your SMB or NFS client.

4. Copy the latest version of the file that you saved in step 1 to Amazon S3 using your SMB or NFS client. Do this through your File Gateway.

Error: ObjectMissing

You can get an `ObjectMissing` error when a writer other than the specified File Gateway deletes the specified file from the S3 bucket. Any subsequent uploads to Amazon S3 or retrievals from Amazon S3 for the object fail.

Your S3 File Gateway can generate a cache report that lists all files in the gateway cache that are currently failing to upload to Amazon S3 due to this error. The information in this report can help you work with Support to resolve issues with your gateway, Amazon S3, or IAM configuration. For more information, see [Create a cache report](#).

To resolve an ObjectMissing error

If the operation that modifies the file is `S3Upload` or `S3GetObject`, do the following:

1. Save the latest copy of the file to the local file system of your SMB or NFS client (you need this file copy in step 3).
2. Delete the file from the File Gateway using your SMB or NFS client.
3. Copy the latest version of the file that you saved in step 1 using your SMB or NFS client. Do this through your File Gateway.

Error: RoleTrustRelationshipInvalid

You get this error when the IAM role for a file share has a misconfigured IAM trust relationship (that is, the IAM role does not trust the Storage Gateway principal named `storagegateway.amazonaws.com`). As a result, the File Gateway would not be able to get the credentials to run any operations on the S3 bucket that backs the file share.

To resolve an RoleTrustRelationshipInvalid error

- Use the IAM console or IAM API to include `storagegateway.amazonaws.com` as a principal that is trusted by your file share's IAM role. For information about IAM role, see [Tutorial: delegate access across AWS accounts using IAM roles](#).

Error: S3AccessDenied

You can get an `S3AccessDenied` error for a file share's Amazon S3 bucket access AWS Identity and Access Management (IAM) role. In this case, the S3 bucket access IAM role that is specified by `roleArn` in the error doesn't allow the operation involved. The operation isn't allowed because of the permissions for the objects in the directory specified by the Amazon S3 prefix.

Your S3 File Gateway can generate a cache report that lists all files in the gateway cache that are currently failing to upload to Amazon S3 due to this error. The information in this report can help you work with Support to resolve issues with your gateway, Amazon S3, or IAM configuration. For more information, see [Create a cache report](#).

To resolve an S3AccessDenied error

- Modify the Amazon S3 access policy that is attached to `roleArn` in the File Gateway health log to allow permissions for the Amazon S3 operation. Make sure that the access policy allows permission for the operation that caused the error. Also, allow permission for the directory specified in the log for `prefix`. For information about Amazon S3 permissions, see [Specifying permissions in a policy](#) in *Amazon Simple Storage Service User Guide*.

These operations can cause an `S3AccessDenied` error to occur:

- `S3HeadObject`
- `S3GetObject`
- `S3ListObjects`
- `S3DeleteObject`
- `S3PutObject`

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: Security scans show open NFS ports

Certain NFS ports are enabled by default, even on gateways that you only use with SMB file shares. If you use third-party security software such as Qualys to scan the network where your File Gateway is deployed, the scan results might report these open NFS ports as a potential security vulnerability. If you only use your gateway with SMB file shares and you want to disable the unused NFS ports for security reasons, use the following procedure:

To disable NFS ports on a File Gateway:

1. Access the gateway local console command prompt using the procedure outlined in <https://docs.aws.amazon.com/filegateway/latest/files3/MaintenanceGatewayConsole-fgw.html>.
2. Enter the following commands to disable NFS traffic:

```
iptables -I INPUT -p udp -m udp --dport 111 -j DROP
iptables -I INPUT -p udp -m udp --dport 2049 -j DROP
iptables -I INPUT -p udp -m udp --dport 20048 -j DROP
iptables -I INPUT -p tcp -m tcp --dport 111 -j DROP
iptables -I INPUT -p tcp -m tcp --dport 2049 -j DROP
iptables -I INPUT -p tcp -m tcp --dport 20048 -j DROP
```

3. Enter the following command to confirm that the blocked NFS ports appear in the IP tables:

```
iptables -n -L -v --line-numbers
```

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](#)
- [Your gateway is slow transferring data to Amazon S3](#)
- [Your gateway is performing more Amazon S3 operations than expected](#)
- [You do not see files in your Amazon S3 bucket](#)
- [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 Amazon S3. 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 S3 bucket 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.

Your gateway is slow transferring data to Amazon S3

If your File Gateway is slow transferring data to Amazon S3, 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 S3. Consider increasing the bandwidth for upload from your File Gateway, adding one or more cache disks, or slowing down client writes.
- 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.
- If `S3PutObjectRequestTime`, `S3UploadPartRequestTime`, or `S3GetObjectRequestTime` are high, there might be a network bottleneck. Try analyzing your network to verify that the gateway has the expected bandwidth.

Your gateway is performing more Amazon S3 operations than expected

If your File Gateway is performing more Amazon S3 operations than expected, check the `FilesRenamed` metric. Rename operations are expensive to run in Amazon S3. Optimize your workflow to minimize the number of rename operations.

You do not see files in your Amazon S3 bucket

If you notice that files on the gateway are not reflected in the Amazon S3 bucket, 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.

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 Amazon S3, 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 Amazon S3, 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.

For NFS, make sure that the client is mounted using a hard mount instead of a soft mount.

Troubleshooting: file share issues

You can find information following about actions to take if you experience unexpected issues with your file share.

Topics

- [Your file share is stuck in CREATING status](#)
- [You can't create a file share](#)
- [SMB file shares don't allow multiple different access methods](#)
- [Multiple file shares can't write to the mapped S3 bucket](#)
- [Notification for deleted log group when using audit logs](#)
- [Can't upload files into your S3 bucket](#)
- [Can't change the default encryption to use SSE-KMS to encrypt objects stored in my S3 bucket](#)
- [Changes made directly in an S3 bucket with object versioning turned on may affect what you see in your file share](#)
- [When writing to an S3 bucket with versioning turned on, the Amazon S3 File Gateway may create multiple versions of Amazon S3 objects](#)
- [Changes to an S3 bucket are not reflected in Storage Gateway](#)
- [ACL permissions aren't working as expected](#)
- [Your gateway performance declined after you performed a recursive operation](#)

Your file share is stuck in CREATING status

When your file share is being created, the status is `CREATING`. The status transitions to `AVAILABLE` status after the file share is created. If your file share gets stuck in the `CREATING` status, do the following:

1. Open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. Make sure the S3 bucket that you mapped your file share to exists. If the bucket doesn't exist, create it. After you create the bucket, the file share status transitions to AVAILABLE. For information about how to create an S3 bucket, see [Create a bucket](#) in the *Amazon Simple Storage Service User Guide*.
3. Make sure your bucket name complies with the rules for bucket naming in Amazon S3. For more information, see [Rules for bucket naming](#) in the *Amazon Simple Storage Service User Guide*.

 **Note**

S3 File Gateway does not support support Amazon S3 buckets with periods (.) in the bucket name.

4. Make sure the IAM role you used to access the S3 bucket has the correct permissions and verify that the S3 bucket is listed as a resource in the IAM policy. For more information, see [Granting access to an Amazon S3 bucket](#).

You can't create a file share

1. If you can't create a file share because your file share is stuck in CREATING status, verify that the S3 bucket you mapped your file share to exists. For information on how to do so, see [Your file share is stuck in CREATING status](#), preceding.
2. If the S3 bucket exists, then verify that AWS Security Token Service is activated in the region where you are creating the file share. If a security token is not active, you should activate it. For information about how to activate a token using AWS Security Token Service, see [Activating and deactivating AWS STS in an AWS Region](#) in the *IAM User Guide*.

SMB file shares don't allow multiple different access methods

SMB file shares have the following restrictions:

1. When the same client attempts to mount both an Active Directory and Guest access SMB file share the following error message is displayed: Multiple connections to a server or shared resource by the same user, using more than one user name, are not allowed. Disconnect all previous connections to the server or shared resource and try again.

2. A Windows user cannot remain connected to two Guest Access SMB file shares, and may be disconnected when a new Guest Access connection is established.
3. A Windows client can't mount both a Guest Access and an Active Directory SMB file share that is exported by the same gateway.

Multiple file shares can't write to the mapped S3 bucket

We don't recommend configuring your S3 bucket to allow multiple file shares to write to one S3 bucket. This approach can cause unpredictable results.

Instead, we recommend that you allow only one file share to write to each S3 bucket. You create a bucket policy to allow only the role associated with your file share to write to the bucket. For more information, see [Best Practices for File Gateway](#).

Notification for deleted log group when using audit logs

If the log group does not exist, the user could select the log group link below that message to go either create a new log group or use an existing log group to use as the target for audit logs

Can't upload files into your S3 bucket

If you can't upload files into your S3 bucket, do the following:

1. Make sure you have granted the required access for the Amazon S3 File Gateway to upload files into your S3 bucket. For more information, see [Granting access to an Amazon S3 bucket](#).
2. Make sure the role that created the bucket has permission to write to the S3 bucket. For more information, see [Best Practices for File Gateway](#).
3. If your File Gateway uses SSE-KMS or DSSE-KMS for encryption, make sure the IAM role associated with the file share includes *kms:Encrypt*, *kms:Decrypt*, *kms:ReEncrypt**, *kms:GenerateDataKey*, and *kms:DescribeKey* permissions. For more information, see [Using Identity-Based Policies \(IAM Policies\) for Storage Gateway](#).

Can't change the default encryption to use SSE-KMS to encrypt objects stored in my S3 bucket

If you change the default encryption and make SSE-KMS (server-side encryption with AWS KMS–managed keys) the default for your S3 bucket, objects that a Amazon S3 File Gateway stores in the

bucket are not encrypted with SSE-KMS. By default, a S3 File Gateway uses server-side encryption managed with Amazon S3 (SSE-S3) when it writes data to an S3 bucket. Changing the default won't automatically change your encryption.

To change the encryption to use SSE-KMS with your own AWS KMS key, you must turn on SSE-KMS encryption. To do so, you provide the Amazon Resource Name (ARN) of the KMS key when you create your file share. You can also update KMS settings for your file share by using the `UpdateNFSFileShare` or `UpdateSMBFileShare` API operation. This update applies to objects stored in the S3 buckets after the update. For more information, see [Data encryption using AWS KMS](#).

Changes made directly in an S3 bucket with object versioning turned on may affect what you see in your file share

If your S3 bucket has objects written to it by another client, your view of the S3 bucket might not be up-to-date as a result of S3 bucket object versioning. You should always refresh your cache before examining files of interest.

Object versioning is an optional S3 bucket feature that helps protect data by storing multiple copies of the same-named object. Each copy has a separate ID value, for example `file1.jpg: ID="xxx"` and `file1.jpg: ID="yyy"`. The number of identically named objects and their lifetimes is controlled by Amazon S3 lifecycle policies. For more details on these Amazon S3 concepts, see [Using versioning](#) and [Object lifecycle management](#) in the *Amazon S3 Developer Guide*.

When you delete a versioned object, that object is flagged with a delete marker but retained. Only an S3 bucket owner can permanently delete an object with versioning turned on.

In your S3 File Gateway, files shown are the most recent versions of objects in an S3 bucket at the time the object was fetched or the cache was refreshed. S3 File Gateways ignore any older versions or any objects marked for deletion. When reading a file, you read data from the latest version. When you write a file in your file share, your S3 File Gateway creates a new version of a named object with your changes, and that version becomes the latest version.

Your S3 File Gateway continues to read from the earlier version, and updates that you make are based on the earlier version should a new version be added to the S3 bucket outside of your application. To read the latest version of an object, use the [RefreshCache](#) API action or refresh from the console as described in [Refreshing Amazon S3 bucket object cache](#).

⚠ Important

We don't recommend that objects or files be written to your S3 File Gateway S3 bucket from outside of the file share.

When writing to an S3 bucket with versioning turned on, the Amazon S3 File Gateway may create multiple versions of Amazon S3 objects

With object versioning turned on, you may have multiple versions of an object created in Amazon S3 on every update to a file from your NFS or SMB client. Here are scenarios that can result in multiple versions of an object being created in your S3 bucket:

- When a file is modified in the Amazon S3 File Gateway by an NFS or SMB client after it has been uploaded to Amazon S3, the S3 File Gateway uploads the new or modified data instead of uploading the whole file. The file modification results in a new version of the Amazon S3 object being created.
- When a file is written to the S3 File Gateway by an NFS or SMB client, the S3 File Gateway uploads the file's data to Amazon S3 followed by its metadata, (ownerships, timestamps, etc.). Uploading the file data creates an Amazon S3 object, and uploading the metadata for the file updates the metadata for the Amazon S3 object. This process creates another version of the object, resulting in two versions of an object.
- When the S3 File Gateway is uploading larger files, it might need to upload smaller chunks of the file before the client is done writing to the File Gateway. Some reasons for this include to free up cache space or a high rate of writes to a file. This can result in multiple versions of an object in the S3 bucket.

You should monitor your S3 bucket to determine how many versions of an object exist before setting up lifecycle policies to move objects to different storage classes. You should configure lifecycle expiration for previous versions to minimize the number of versions you have for an object in your S3 bucket. The use of Same-Region replication (SRR) or Cross-Region replication (CRR) between S3 buckets will increase the storage used. For more information about replication, see [Replication](#).

⚠ Important

Do not configure replication between S3 buckets until you understand how much storage is being used when object versioning is turned on.

Use of versioned S3 buckets can greatly increase the amount of storage in Amazon S3 because each modification to a file creates a new version of the S3 object. By default, Amazon S3 continues to store all of these versions unless you specifically create a policy to override this behavior and limit the number of versions that are kept. If you notice unusually large storage usage with object versioning turned on, check that you have your storage policies set appropriately. An increase in the number of HTTP 503-slow down responses for browser requests can also be the result of problems with object versioning.

If you turn on object versioning after installing a S3 File Gateway, all unique objects are retained (ID="NULL") and you can see them all in the file system. New versions of objects are assigned a unique ID (older versions are retained). Based on the object's timestamp only the newest versioned object is viewable in the NFS file system.

After you turn on object versioning, your S3 bucket can't be returned to a nonversioned state. You can, however, suspend versioning. When you suspend versioning, a new object is assigned an ID. If the same named object exists with an ID="NULL" value, the older version is overwritten. However, any version that contains a non-NULL ID is retained. Timestamps identify the new object as the current one, and that is the one that appears in the NFS file system.

Changes to an S3 bucket are not reflected in Storage Gateway

Storage Gateway updates the file share cache automatically when you write files to the cache locally using the file share. However, Storage Gateway doesn't automatically update the cache when you upload a file directly to Amazon S3. When you do this, you must perform a `RefreshCache` operation to see the changes on the file share. If you have more than one file share, then you must run the `RefreshCache` operation on each file share.

You can refresh the cache using the Storage Gateway console and the AWS Command Line Interface (AWS CLI):

- To refresh the cache using the Storage Gateway console, see [Refreshing objects in your Amazon S3 bucket](#).

- To refresh the cache using the AWS CLI:
 1. Run the command `aws storagegateway list-file-shares`
 2. Copy the Amazon Resource Number (ARN) of the file share with the cache that you want to refresh.
 3. Run the `refresh-cache` command with your ARN as the value for `--file-share-arn`:

```
aws storagegateway refresh-cache --file-share-arn
arn:aws:storagegateway:eu-west-1:12345678910:share/share-FFDEE12
```

To automate the `RefreshCache` operation, see [How can I automate the RefreshCache operation on Storage Gateway?](#)

ACL permissions aren't working as expected

If access control list (ACL) permissions aren't working as you expect with your SMB file share, you can perform a test.

To do this, first test the permissions on a Microsoft Windows file server or a local Windows file share. Then compare the behavior to your gateway's file share.

Your gateway performance declined after you performed a recursive operation

In some cases, you might perform a recursive operation, such as renaming a directory or turning on inheritance for an ACL, and force it down the tree. If you do this, your S3 File Gateway recursively applies the operation to all objects in the file share.

For example, suppose that you apply inheritance to existing objects in an S3 bucket. Your S3 File Gateway recursively applies inheritance to all objects in the bucket. Such operations can cause your gateway performance to decline.

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](#)
- [Metrics](#)

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](#)
- [Notification: HardReboot](#)
- [Notification: HealthCheckFailure](#)
- [Notification: AvailabilityMonitorTest](#)

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](#)
- [Best practices: managing multipart uploads](#)
- [Best practices: Unzip compressed files locally before copying to a gateway](#)
- [Retain file attributes when copying data from Windows Server](#)
- [Best practices: Proper sizing of cache disks](#)
- [Working with multiple file shares and Amazon S3 buckets](#)
- [Clean up unnecessary resources](#)

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 to the Amazon S3 bucket 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 S3 File Gateway](#).
2. Create a new on the EC2 gateway you created. For more information, see [Create a file share](#).
3. Mount your file share on your client and map it to the S3 bucket that contains the data that you want to recover. For more information, see [Mount and use your file share](#).

Best practices: managing multipart uploads

When transferring large files, S3 File Gateway makes use of the Amazon S3 multipart upload feature to split the files into smaller parts and transfer them in parallel for improved efficiency. For more information about multipart upload, see [Uploading and copying objects using multipart upload](#) in the *Amazon Simple Storage Service User Guide*.

If a multipart upload doesn't complete successfully for any reason, the gateway typically stops the transfer, deletes any partially-transferred pieces of the file from Amazon S3, and attempts the transfer again. In rare cases, such as when hardware or network failure prevent the gateway from cleaning up after an unsuccessful multipart upload, pieces of the partially-transferred file might remain on Amazon S3 where they can incur storage charges.

As a best practice for minimizing Amazon S3 storage costs from incomplete multipart uploads, we recommend configuring an Amazon S3 bucket lifecycle rule that uses the `AbortIncompleteMultipartUpload` API action to automatically stop unsuccessful transfers and delete associated file parts after a designated number of days. For instructions, see [Configuring a bucket lifecycle configuration to delete incomplete multipart uploads](#) in the *Amazon Simple Storage Service User Guide*.

Best practices: Unzip compressed files locally before copying to a gateway

If you try to unzip a compressed archive containing thousands of files while it is stored on your gateway, you might encounter significant performance-related delays. The process of unzipping an archive that contains large numbers of files on any type of network file share inherently involves a high volume of input/output operations, metadata cache manipulation, network overhead, and latency. Additionally, Storage Gateway is unable to determine when each file from the archive has finished unzipping, and can begin uploading files before the process is complete, which further impacts performance. These issues are compounded when the files inside the archive are numerous, but small in size.

As a best practice, we recommend transferring compressed archives from your gateway to your local machine first, before you unzip them. Then, if necessary, you can use a tool such as *robocopy* or *rsync* to transfer the unzipped files back to the gateway.

Retain file attributes when copying data from Windows Server

It is possible to copy files to your File Gateway using the basic copy command on Microsoft Windows, but this command copies only the file data by default - omitting certain file attributes such as security descriptors. If the files are copied to the gateway without the corresponding security restrictions and Discretionary Access Control List (DACL) information, it is possible that they could be accessed by unauthorized users.

As a best practice for preserving all file attributes and security information when copying files to your gateway on Microsoft Windows Server, we recommend using the `robocopy` or `xcopy` commands, with the `/copy:DS` or `/o` flags, respectively. For more information, see [robocopy](#) and [xcopy](#) in the Microsoft Windows Server command reference documentation.

Best practices: Proper sizing of cache disks

For best performance, the total disk cache size must be large enough to cover the size of your active working set. For read-heavy and mixed read/write workloads, this ensures that you can achieve a high percentage of cache hits on reads, which is desirable. You can monitor this via the `CacheHitPercent` metric for your S3 File Gateway.

For write-heavy workloads (e.g. for backup and archival), the S3 File Gateway buffers incoming writes on the disk cache prior to copying this data asynchronously to Amazon S3. You should ensure that you have sufficient cache capacity to buffer written data. The `CachePercentDirty` metric provides an indication of the percentage of the disk cache that has not yet been persisted to AWS.

Low values of `CachePercentDirty` are desirable. Values that are consistently close to 100% indicate that the S3 File Gateway is unable to keep up with the rate of incoming write traffic. You can avoid this by either increasing the provisioned disk cache capacity, or increasing the dedicated network bandwidth available from the S3 File Gateway to Amazon S3, or both.

For more information about cache disk sizing, see [Amazon S3 File Gateway cache sizing best practices](#) on the official Amazon Web Services YouTube channel.

Working with multiple file shares and Amazon S3 buckets

When you configure a single Amazon S3 bucket to allow multiple gateways or file shares to write to it, the results can be unpredictable. You can configure your buckets in one of two ways to avoid

unpredictable results. Choose the configuration method that best fits your use case from the following options:

- Configure your S3 buckets so that only one file share can write to each bucket. Use a different file share to write to each bucket.

To do this, create an S3 bucket policy that denies all roles except for the role that's used for a specific file share to put or delete objects in the bucket. Attach a similar policy to each bucket, specifying a different file share to write to each bucket.

The following example policy denies S3 bucket write permissions to all roles except for the role that created the bucket. The `s3:DeleteObject` and `s3:PutObject` actions are denied for all roles except "TestUser". The policy applies to all objects in the "arn:aws:s3:::amzn-s3-demo-bucket/*" bucket.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DenyMultiWrite",
      "Effect": "Deny",
      "Principal": "*",
      "Action": [
        "s3:DeleteObject",
        "s3:PutObject"
      ],
      "Resource": "arn:aws:s3:::amzn-s3-demo-bucket/*",
      "Condition": {
        "StringNotLike": {
          "aws:userid": "TestUser:*"
        }
      }
    }
  ]
}
```

- If you do want multiple file shares to write to the same Amazon S3 bucket, you must prevent the file shares from trying to write to the same objects simultaneously.

To do this, configure a separate, unique object prefix for each file share. This means that each file share only writes to objects with the corresponding prefix, and doesn't write to objects that are associated with the other file shares in your deployment. You configure the object prefix in the **S3 prefix name** field when you create a new file share.

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.
- [File attribute support](#) - Learn how your gateway handles DOS and Windows file attributes.
- [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.
- [Using storage classes](#) - Learn about the Amazon S3 storage classes that File Gateway supports, and what to consider when choosing a storage class.
- [Using Kubernetes CSI drivers](#) - Learn how to install and configure Container Storage Interface (CSI) drivers to allow Kubernetes instances to use File Gateway for storage.
- [Terraform module](#) - Learn how to use Terraform to deploy File Gateway as a virtual machine.

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 S3 File Gateway](#)
- [Deploy a customized Amazon EC2 host for S3 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 S3 File Gateway

This topic lists the steps to deploy an Amazon EC2 host using the default specifications.

You can deploy and activate an Amazon S3 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 S3 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				
111	TCP, UDP	NFSv3				
139	TCP, UDP	SMB				
445	TCP	SMB				
2049	TCP, UDP	NFS				
20048	TCP, UDP	NFSv3				

- **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				

Default Settings	AMI Root Volume	Volume 2 Cache				
Throughput	125	125				
t						

Deploy a customized Amazon EC2 host for S3 File Gateway

You can deploy and activate an Amazon S3 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.

S3 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_S3-1.25.0`

To deploy an Amazon EC2 instance to host your Amazon S3 File Gateway

1. Start setting up a new gateway using the Storage Gateway console. For instructions, see [Set up an Amazon S3 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 S3 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 S3 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. If you plan to create NFS file shares, you must open TCP/UDP port 2049 for NFS access, TCP/UDP port 111 for NFSv3 access, and TCP/UDP port 20048 for NFSv3 access. If you plan to create SMB file shares, you must open TCP port 445 for SMB 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 S3 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 S3 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_S3/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_S3/latest
```

The CLI command returns output similar to the following.

```
{
  "Parameter": {
    "Type": "String",
    "LastModifiedDate": 1561054105.083,
    "Version": 4,
    "ARN": "arn:aws:ssm:us-east-2::parameter/aws/service/storagegateway/ami/FILE_S3/latest",
    "Name": "/aws/service/storagegateway/ami/FILE_S3/latest",
    "Value": "ami-123c45dd67d891000"
  }
}
```

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](#)
- [Configuring Your Gateway for Multiple NICs in Microsoft Hyper-V Host](#)

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.

 **Note**

The example mounting commands provided on the info page for a file share in the Storage Gateway console will always include the IP address of the network adapter that was most recently added to the file share's associated gateway.

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.

 **Note**

The example mounting commands provided on the info page for a file share in the Storage Gateway console will always include the IP address of the network adapter that was most recently added to the file share's associated gateway.

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 S3 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 S3 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 S3 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 \(curl\)](#)
- [Linux \(bash/zsh\)](#)
- [Microsoft Windows PowerShell](#)
- [Using your local console](#)

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.

Support for file attributes in Amazon S3 File Gateway

Amazon S3 File Gateway supports DOS or Windows file attributes by default. Using S3 File Gateway, you can preserve file data and metadata and update settings — such as marking items as archived when they are placed in Amazon S3. For more information about DOS and Windows file attributes, see the [File Attribute Constants](#) article on the Windows app development documentation website.

S3 File Gateway supports the following attributes:

- *ReadOnly* – The S3 File Gateway prevents changes to files that have the ReadOnly attribute set.
- *Archive* – The S3 File Gateway sets this attribute when files are first added to the gateway.

Note

Backup applications commonly backup files that have the Archive bit set and then clear the bit after successful backup.

- *Hidden* – Server Message Block (SMB) clients hide files that use this bit set.
- *System* – This attribute persists once you have set it.

When you copy a file to the S3 File Gateway with the attributes set, the file's DOS or Windows attributes are preserved on the S3 File Gateway and in Amazon S3. You can update these attributes for files on the gateway, and those updates also apply to the object in Amazon S3. If a file is evicted from the gateway the gateway pull the file, its metadata, and its persistent attributes from Amazon S3 when you request.

Note

DOS attributes are only supported on SMB shares and if access is controlled by Windows Access Control Lists.

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 shares 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: <i>region</i> : <i>account-id</i> :gateway/ <i>gateway-id</i>
File Share ARN	arn:aws:storagegateway: <i>region</i> : <i>account-id</i> :share/ <i>share-id</i>

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](#)
- [Open-source components for Amazon S3 File Gateway](#)

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 S3 File Gateway

Several third-party tools and licenses are used to deliver Amazon S3 File Gateway (S3 File Gateway) functionality.

Use the following links to download the source code for certain open-source software components that are included with S3 File Gateway software:

- For Amazon S3 File Gateway: [sgw-file-s3-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 [Third-Party Licenses](#).



Limits and quotas for Amazon S3 File Gateway



Quotas for file shares

The following table lists quotas for file shares.

Description	Limit
Maximum number of file shares per gateway	50

Description	Limit
<p> Note</p> <p>Each file share can only connect to one S3 bucket, but multiple file shares can connect to the same bucket. If you connect more than one file share to the same bucket, you must configure each file share to use a unique, non-overlapping prefix name to prevent read/write conflicts.</p> <p>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.</p>	

Description	Limit
<p>The maximum number of files for which a gateway can simultaneously cache metadata</p> <div data-bbox="115 352 792 1287"><p> Note</p><p>Because S3 File Gateway is backed by Amazon S3, there is no maximum folder size or limit to the number of files that you can store or access using a gateway.</p><p>Each gateway has a configurable limit that determines the number of files for which it can simultaneously cache metadata. You can use the UpdateGatewayInformation API action to set GatewayCapacity to Small, Medium, or Large. This setting impacts gateway performance and hardware recommendations. For more information, see Performance guidance for gateways with multiple file shares.</p></div>	<p>Gateway Capacity:</p> <p>Small — 5M files</p> <p>Medium — 10M files</p> <p>Large — 20M files</p>
<p>The maximum size of an individual file</p> <div data-bbox="115 1402 792 1862"><p> Note</p><p>If you try to write files larger than the size limit piece by piece, only the first 5 TiB will be uploaded. If you try to write the files larger than the size limit all at once, on Windows clients no file is created and on Linux clients a file of size zero is created.</p></div>	<p>5 TiB</p>

Description	Limit
Maximum path length	1024 bytes
<div><div> Note</div><div>Clients are not allowed to create a path exceeding this length, and doing so results in an error. This limit applies to both protocols supported by File Gateways, NFS and SMB. Path length in bytes is calculated from character bit values in UTF-8 encoding.</div></div>	
Maximum file name length	255 bytes
<div><div> Note</div><div>File Gateway does not support file names that exceed this length. File name length in bytes is calculated from character bit values in UTF-8 encoding.</div></div>	

Recommended local disk sizes for your gateway

The following table recommends sizes for local disk storage for your deployed gateway.

Gateway Type	Cache (Minimum)	Cache (Maximum)	Other Required Local Disks
S3 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.

Using storage classes

Amazon S3 File Gateway supports the Amazon S3 Standard, Amazon S3 Standard-Infrequent Access, Amazon S3 One Zone-Infrequent Access, Amazon S3 Intelligent-Tiering, and S3 Glacier storage classes. For more information about storage classes, see [Amazon S3 storage classes](#) in the *Amazon Simple Storage Service User Guide*.

Note

S3 File Gateway does not currently support the Amazon S3 Glacier Instant Retrieval storage class.

Topics

- [Using storage classes with a File Gateway](#)
- [Using the GLACIER storage class with File Gateway](#)

Using storage classes with a File Gateway

When you create or update a file share, you have the option to select a storage class for your objects. You can choose the Amazon S3 Standard storage class, or any of the S3 Standard-IA, S3 One Zone-IA, or S3 Intelligent-Tiering storage classes. Objects stored in any of these storage classes can be transitioned to GLACIER using a lifecycle policy.

Amazon S3 storage class	Considerations
Standard	Choose Standard to store your frequently accessed files redundantly in multiple Availabil

Amazon S3 storage class	Considerations
	<p>ity Zones that are geographically separated. This is the default storage class. See Amazon S3 pricing for more details.</p>

Amazon S3 storage class	Considerations
S3 Intelligent-Tiering	<p>Choose Intelligent-Tiering to optimize storage costs by automatically moving data to the most cost-effective storage access tier.</p> <p>Objects smaller than 128 KB are not eligible for auto tiering in the Intelligent-Tiering storage class. These objects are charged at the frequent access tier rates, and don't incur the monitoring fee charged for auto-tiered objects.</p> <p>S3 Intelligent-Tiering now supports an Archive Access tier and a Deep Archive Access tier. S3 Intelligent-Tiering automatically moves objects that haven't been accessed for 90 days to the Archive Access tier, and after 180 days without being accessed, to the Deep Archive Access tier. Whenever an object in one of the archive access tiers is restored, the object moves to the Frequent Access tier within a few hours and is ready to be retrieved. This creates timeout errors for users or applications trying to access files through a file share if the object only exists in one of the two archive tiers. Don't use the archive tiers with S3 Intelligent-Tiering if your applications are accessing files through the file shares that are presented by the File Gateway.</p> <p>When file operations that update metadata (such as owner, timestamp, permissions, and ACLs) are performed against files managed by the File Gateway, the existing object is deleted and a new version of the object is created in this Amazon S3 storage class. Validate how file</p>

Amazon S3 storage class	Considerations
	<p>operations impact object creation before using this storage class in production. See Amazon S3 pricing for more details.</p>
S3 Standard-IA	<p>Choose Standard-IA to store your infrequently accessed files redundantly in multiple Availability Zones that are geographically separated.</p> <p>Objects stored in the Standard-IA storage class can incur additional charges for overwriting, deleting, requesting, retrieving, or transitioning objects between storage classes within 30 days. There is a minimum storage duration of 30 days. Objects deleted before 30 days incur a pro-rated charge equal to the storage charge for the remaining days. Consider how often these objects change, how long you plan to keep these objects, and how often you need to access them. Objects smaller than 128 KB are charged for 128 KB and early deletion fees apply.</p> <p>When file operations that update metadata (such as owner, timestamp, permissions, and ACLs) are performed against files managed by the File Gateway, the existing object is deleted and a new version of the object is created in this Amazon S3 storage class. You should validate how file operations impact object creation before using this storage class in production because early deletion fees apply. See Amazon S3 pricing for more details.</p>

Amazon S3 storage class	Considerations
S3 One Zone-IA	<p>Choose One Zone-IA to store your infrequently accessed files in a single Availability Zone.</p> <p>Objects stored in the One Zone-IA storage class can incur additional charges for overwriting, deleting, requesting, retrieving, or transitioning objects between storage classes within 30 days. There is a minimum storage duration of 30 days, and objects deleted before 30 days incur a pro-rated charge equal to the storage charge for the remaining days. Consider how often these objects change, how long you plan to keep these objects, and how often you need to access them. Objects smaller than 128 KB are charged for 128 KB and early deletion fees apply.</p> <p>When file operations that update metadata (such as owner, timestamp, permissions, and ACLs) are performed against files managed by the File Gateway, the existing object is deleted and a new version of the object is created in this Amazon S3 storage class. You should validate how file operations impact object creation before using this storage class in production because early deletion fees apply. See Amazon S3 pricing for more details.</p>

Although you can write objects directly from a file share to the S3-Standard-IA, S3-One Zone-IA, or S3 Intelligent-Tiering storage class, we recommend that you use a lifecycle policy to transition your objects rather than write directly from the file share, especially if you're expecting to update or delete the object within 30 days of archiving it. For information about lifecycle policy, see [Object lifecycle management](#).

Using the GLACIER storage class with File Gateway

If you transition a file to S3 Glacier through Amazon S3 lifecycle policies, and the file is visible to your file share clients through the cache, you will encounter I/O errors when you update the file. We recommend that you set up CloudWatch Events to receive notification when these I/O errors occur, and use the notification to take action. For example, you can take action to restore the archived object to Amazon S3. After the object is restored to S3, your file share clients can access and update them successfully through the file share.

For information about how to restore archived objects, see [Restoring archived objects](#) in the *Amazon Simple Storage Service User Guide*.

Important

S3 File Gateway does not officially support the S3 Glacier Instant Retrieval storage class. Although you can designate objects in a file share bucket for S3 Glacier Instant Retrieval by using lifecycle policies or direct PUT requests, S3 File Gateway cannot recognize which files are in that storage class, and will perform file operations on them like any other object. Because S3 Glacier Instant Retrieval has higher costs for access than other Amazon S3 storage classes, bulk file operations such as virus scans, rsync, and renames, can result in large Amazon S3 bills if not managed carefully. For this reason, we do not recommend using S3 Glacier Instant Retrieval with S3 File Gateway.

Using Kubernetes Container Storage Interface drivers

Kubernetes is an open-source system for automating deployment, scaling, and management of containerized applications. In a Kubernetes environment, a container is similar to a VM, but a container has relaxed isolation properties to share the Operating System (OS) among its applications. Therefore, containers are considered to be more lightweight than VMs. Similar to a VM, a container has its own filesystem, a share of allocated CPU, memory, process space, and more. As they are decoupled from the underlying infrastructure, they are portable across clouds and OS distributions. If you have a Kubernetes cluster, you can install and configure Kubernetes Container Storage Interface (CSI) drivers across the instances in your cluster to allow them to use an existing Amazon S3 File Gateway for storage.

After you install the CSI drivers for the type of file share that you want to use, you must create one or more storage objects. Depending on the type of provisioning that you want Kubernetes to

use when your pods request storage, you must create either a single Kubernetes `StorageClass` object, or both a `PersistentVolume` object *and* a `PersistentVolumeClaim` object to connect your Kubernetes compute pods to your file share. For more information, refer to the Kubernetes online documentation at <https://kubernetes.io/docs/concepts/storage/>.

Topics

- [Working with SMB CSI drivers](#)
- [Working with NFS CSI drivers](#)

Working with SMB CSI drivers

Follow the procedures in this section to install, configure, or delete the CSI drivers that are required to use an SMB file share on an Amazon S3 File Gateway for storage in your Kubernetes cluster. For more information, see the open-source SMB CSI driver documentation on GitHub at <https://github.com/kubernetes-csi/csi-driver-smb/blob/master/docs/install-csi-driver-master.md>.

Note

When you create a `PersistentVolume` object or a `StorageClass` object, you can specify a `ReclaimPolicy` parameter to determine what happens to the external storage when objects are deleted. The SMB CSI driver supports the `Retain` and `Recycle` options, but does not currently support a `Delete` option.

Install drivers

To install Kubernetes SMB CSI drivers:

1. From a command-line terminal with access to `kubectl` for your Kubernetes cluster, run the following command:

```
curl -skSL https://raw.githubusercontent.com/kubernetes-csi/csi-driver-smb/master/
deploy/install-driver.sh | bash -s master --
```

2. Wait for the previous command to finish, then use the following commands to ensure that the CSI driver pods are running:

```
kubectl -n kube-system get pod -o wide --watch -l app=csi-smb-controller
```

```
kubectl -n kube-system get pod -o wide --watch -l app=csi-smb-node
```

The output should look similar to the following:

NAME	READY	STATUS	RESTARTS	AGE	IP
csi-smb-controller-56bfddd689-dh5tk	4/4	Running	0	35s	
10.240.0.19 k8s-agentpool-22533604-0					
csi-smb-controller-56bfddd689-8pgr4	4/4	Running	0	35s	
10.240.0.35 k8s-agentpool-22533604-1					
csi-smb-node-cvgbs	3/3	Running	0	35s	
10.240.0.35 k8s-agentpool-22533604-1					
csi-smb-node-dr4s4	3/3	Running	0	35s	
10.240.0.4 k8s-agentpool-22533604-0					

Create an SMB StorageClass object

To create a new SMB StorageClass object for your Kubernetes cluster:

1. Create a configuration file named `storageclass.yaml` with contents similar to the following example. Substitute your own deployment-specific information for the *ExampleValues* shown.

```
---
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: ExampleStorageClassName
provisioner: smb.csi.k8s.io
parameters:
  source: "//gateway-dns-name-or-ip-address/example-share-name"
  # if csi.storage.k8s.io/provisioner-secret is provided, will create a sub
  directory
  # with PV name under source
  csi.storage.k8s.io/provisioner-secret-name: "examplesmbcreds"
  csi.storage.k8s.io/provisioner-secret-namespace: "examplnamespace"
  csi.storage.k8s.io/node-stage-secret-name: "examplesmbcreds"
  csi.storage.k8s.io/node-stage-secret-namespace: "examplnamespace"
```

```
volumeBindingMode: Immediate
reclaimPolicy: Retain
mountOptions:
  - dir_mode=0777
  - file_mode=0777
  - uid=1001
  - gid=1001
```

2. From a command-line terminal with access to `kubectl` and `storageclass.yaml`, run the following command:

`kubectl apply -f storageclass.yaml`

Note

You can also create the `StorageClass` by providing the `.yaml` configuration text from the previous step to most third-party Kubernetes management and containerization platforms.

3. Configure the pods in your Kubernetes cluster to use the new `StorageClass` that you created. For more information, refer to the Kubernetes online documentation at <https://kubernetes.io/docs/concepts/storage/>.

Create SMB `PersistentVolume` and `PersistentVolumeClaim` objects

To create new `SMB PersistentVolume` and `PersistentVolumeClaim` objects:

1. Create two configuration files. One named `persistentvolume.yaml`, and one named `persistentvolumeclaim.yaml`.
2. For `persistentvolume.yaml`, add contents that are similar to the following example. Substitute your own deployment-specific information for the *ExampleValues* shown.

```
---
apiVersion: v1
kind: PersistentVolume
metadata:
  name: pv-smb-example-name
  namespace: smb-example-namespace # PersistentVolume and PersistentVolumeClaim
  must use the same namespace parameter
```

```
spec:
  capacity:
    storage: 100Gi
  accessModes:
    - ReadWriteMany
  persistentVolumeReclaimPolicy: Retain
  mountOptions:
    - dir_mode=0777
    - file_mode=0777
    - vers=3.0
  csi:
    driver: smb.csi.k8s.io
    readOnly: false
    volumeHandle: examplehandle # make sure it's a unique id in the cluster
    volumeAttributes:
      source: "//gateway-dns-name-or-ip-address/example-share-name"
    nodeStageSecretRef:
      name: example-smbcreds
      namespace: smb-example-namespace
```

3. For `persistentvolumeclaim.yaml`, add contents that are similar to the following example. Substitute your own deployment-specific information for the *ExampleValues* shown.

```
---
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: examplename-pvc-smb-static
  namespace: smb-example-namespace # PersistentVolume and PersistentVolumeClaim
  must use the same namespace parameter
spec:
  accessModes:
    - ReadWriteMany
  resources:
    requests:
      storage: 10Gi
      volumeName: pv-smb-example-name # make sure specified volumeName matches the
      name of the PersistentVolume you created
      storageClassName: ""
```

4. From a command-line terminal with access to `kubectl` and both of the `.yaml` files that you created, run the following commands:

```
kubectl apply -f persistentvolume.yaml
```

```
kubectl apply -f persistentvolumeclaim.yaml
```

Note

You can also create the PersistentVolume and PersistentVolumeClaim objects by providing the .yaml configuration text from the previous step to most third-party Kubernetes management and containerization platforms.

5. Configure the pods in your Kubernetes cluster to use the new PersistentVolumeClaim that you created. For more information, refer to the Kubernetes online documentation at <https://kubernetes.io/docs/concepts/storage/>.

Uninstall drivers

To uninstall the Kubernetes SMB CSI drivers:

- From a command-line terminal with access to `kubectl` for your Kubernetes cluster, run the following command:

```
curl -sSL https://raw.githubusercontent.com/kubernetes-csi/csi-driver-smb/master/deploy/uninstall-driver.sh | bash -s --
```

Working with NFS CSI drivers

Follow the procedures in this section to install, configure, or delete the CSI drivers that are required to use an NFS file share on an Amazon S3 File Gateway for storage in your Kubernetes cluster.

For more information, see the open-source NFS CSI driver documentation on GitHub at <https://github.com/kubernetes-csi/csi-driver-nfs/blob/master/docs/install-csi-driver-master.md>.

Install drivers

To install Kubernetes NFS CSI drivers:

1. From a command-line terminal with access to `kubectl` for your Kubernetes cluster, run the following command:

```
curl -skSL https://raw.githubusercontent.com/kubernetes-csi/csi-driver-nfs/master/
deploy/install-driver.sh | bash -s master --
```

2. Wait for the previous command to finish, then use the following commands to ensure that the CSI driver pods are running:

```
kubectl -n kube-system get pod -o wide -l app=csi-nfs-controller
```

```
kubectl -n kube-system get pod -o wide -l app=csi-nfs-node
```

The output should look similar to the following:

NAME	READY	STATUS	RESTARTS	AGE	IP
csi-nfs-controller-56bfddd689-dh5tk	4/4	Running	0	35s	
10.240.0.19 k8s-agentpool-22533604-0					
csi-nfs-controller-56bfddd689-8pgr4	4/4	Running	0	35s	
10.240.0.35 k8s-agentpool-22533604-1					
csi-nfs-node-cvqbs	3/3	Running	0	35s	
10.240.0.35 k8s-agentpool-22533604-1					
csi-nfs-node-dr4s4	3/3	Running	0	35s	
10.240.0.4 k8s-agentpool-22533604-0					

Create an NFS StorageClass object

To create an NFS StorageClass object for your Kubernetes cluster:

1. Create a configuration file named `storageclass.yaml` with contents that are similar to the following example. Substitute your own deployment-specific information for the *ExampleValues* shown.

```
---
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: example-nfs-classname
  namespace: example-namespace
provisioner: nfs.csi.k8s.io
parameters:
```

```
server: gateway-dns-name-or-ip-address
share: /example-share-name
reclaimPolicy: Retain
volumeBindingMode: Immediate
mountOptions:
  - hard
  - nfsvers=4.1
```

2. From a command-line terminal with access to `kubectl` and `storageclass.yaml`, run the following command:

`kubectl apply -f storageclass.yaml`

Note

You can also create the StorageClass by providing the `.yaml` configuration text from the previous step to most third-party Kubernetes management and containerization platforms.

3. Configure the pods in your Kubernetes cluster to use the new StorageClass object that you created. For more information, refer to the Kubernetes online documentation at <https://kubernetes.io/docs/concepts/storage/>.

Create NFS PersistentVolume and PersistentVolumeClaim objects

To create new NFS PersistentVolume and PersistentVolumeClaim objects:

1. Create two configuration files named `persistentvolume.yaml` and `persistentvolumeclaim.yaml`.
2. For `persistentvolume.yaml`, add contents that are similar to the following example. Substitute your own deployment-specific information for the *ExampleValues* shown.

```
---
apiVersion: v1
kind: PersistentVolume
metadata:
  name: pvc-nfs-exemplename
spec:
  capacity:
```

```

    storage: 10Gi
  accessModes:
    - ReadWriteMany
  persistentVolumeReclaimPolicy: Retain
  mountOptions:
    - hard
    - nolock
    - nfsvers=4.1
  csi:
    driver: nfs.csi.k8s.io
    readOnly: false
    volumeHandle: unique-volumeid-example # make sure it's a unique id in the
cluster
    volumeAttributes:
      server: gateway-dns-name-or-ip-address
      share: /example-share-name

```

3. For `persistentvolumeclaim.yaml`, add contents that are similar to the following example. Substitute your own deployment-specific information for the *ExampleValues* shown.

```

---
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: exemplename-pvc-nfs-static
spec:
  accessModes:
    - ReadWriteMany
  resources:
    requests:
      storage: 10Gi
  volumeName: pv-nfs-exemplename # make sure specified volumeName matches the name
of the PersistentVolume you created
  storageClassName: ""

```

4. From a command-line terminal with access to `kubectl` and both `.yaml` files, run the following commands:

```
kubectl apply -f persistentvolume.yaml
```

```
kubectl apply -f persistentvolumeclaim.yaml
```

Note

You can also create the PersistentVolume and PersistentVolumeClaim objects by providing the `.yaml` configuration text from the previous step to most third-party Kubernetes management and containerization platforms.

5. Configure the pods in your Kubernetes cluster to use the new PersistentVolumeClaim object that you created. For more information, refer to the Kubernetes online documentation at <https://kubernetes.io/docs/concepts/storage/>.

Uninstall drivers

To uninstall Kubernetes NFS CSI drivers:

- From a command-line terminal with access to `kubectl` for your Kubernetes cluster, run the following command:

```
curl -skSL https://raw.githubusercontent.com/kubernetes-csi/csi-driver-nfs/master/
deploy/uninstall-driver.sh | bash -s master --
```

AWS Storage Gateway Terraform module

[HashiCorpTerraform](#) is an open-source Infrastructure as Code (IaC) engine developed using the HashiCorp Configuration Language (HCL). Terraform provides a consistent command line interface (CLI) workflow that, in conjunction with Amazon S3 File Gateway for the back-end infrastructure, can manage hundreds of cloud services and codify cloud APIs into declarative configuration files.

You can use Terraform to safely deploy an Amazon S3 File Gateway as a virtual machine (VM) in your on-premises virtual infrastructure. Terraform provides automation for on-premises virtual infrastructure. See [Automate Amazon S3 File Gateway deployments in VMware with Terraform by HashiCorp](#) for information about quickly deploying an Amazon S3 File Gateway using Terraform within an on-premises VMware virtual environment.

Note

You might need to configure Terraform to obtain the latest version of the AWS Storage Gateway machine image for your preferred hypervisor platform. Storage Gateway machine

images use the following naming convention. The version number appended to the image name changes with each version release.

`aws-storage-gateway-FILE_S3-1.25.0`

This automation provides you with a customizable Terraform module that you can use to provision an Amazon S3 File Gateway with all of the resources and dependencies needed to fully deploy the gateway and file shares in your VM environment. The Terraform module provisions the gateway VM, activates the gateway, configures the cache disk, joins the gateway to a domain, creates the Amazon S3 buckets, creates the file shares, and maps them to buckets. For a complete example of a repository that contains Terraform code to create the resources required to run Amazon S3 File Gateway on premises, see the [Terraform Storage Gateway module](#) source code on GitHub.

 **Note**

The Amazon S3 File Gateway module for Terraform is a community supported effort. It is not part of an AWS service. Best-effort support is provided by the AWS Storage community.

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](#)
- [Signing Requests](#)
- [Error Responses](#)
- [Storage Gateway API Actions](#)

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	<p>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):</p> <pre>Authorization: AWS4-HMAC_SHA456 Credentials= <i>YourAccessKey</i> /<i>yyyymmdd</i>/<i>region</i>/storagegateway/aw s4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-targ et, Signature= <i>CalculatedSignature</i></pre> <p>In the preceding syntax, you specify <i>YourAccessKey</i>, the year, month, and day (<i>yyyymmdd</i>), the <i>region</i>, and the <i>CalculatedSignature</i>. 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.</p>
Content-Type	<p>Use application/x-amz-json-1.1 as the content type for all requests to AWS Storage Gateway.</p> <pre>Content-Type: application/x-amz-json-1.1</pre>
Host	<p>Use the host header to specify the AWS Storage Gateway endpoint where you send your request. For example, storagegateway.us-</p>

Header	Description
	<p>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 <i>AWS General Reference</i>.</p> <div data-bbox="472 457 1507 537"> Host: storagegateway. <i>region</i>.amazonaws.com </div>
x-amz-date	<p>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.</p> <div data-bbox="472 1018 1507 1098"> x-amz-date: <i>YYYYMMDD'T'HHMMSS'Z'</i> </div>
x-amz-target	<p>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.</p> <div data-bbox="472 1333 1507 1413"> x-amz-target: StorageGateway_ <i>APIVersion</i> .<i>operationName</i> </div> <p>The <i>operationName</i> value (e.g. "ActivateGateway") can be found from the API list, API Reference for Storage Gateway.</p>

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
92c0effa6f9224ac752ca179a04cecbde3038b0959666a8160ab452c9e51b3e
```

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
<code>IncompleteSignatureException</code>	The specified signature is incomplete.	400 Bad Request
<code>InternalFailure</code>	The request processing has failed due to some unknown error, exception or failure.	500 Internal Server Error
<code>InternalServerError</code>	One of the operation error code messages Operation Error Codes .	500 Internal Server Error
<code>InvalidAction</code>	The requested action or operation is invalid.	400 Bad Request
<code>InvalidClientTokenId</code>	The X.509 certificate or AWS Access Key ID provided does not exist in our records.	403 Forbidden
<code>InvalidGatewayRequestException</code>	One of the operation error code messages in Operation Error Codes .	400 Bad Request
<code>InvalidSignatureException</code>	The request signature we calculated does not match the signature you provided. Check your AWS Access Key and signing method.	400 Bad Request
<code>MissingAction</code>	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
<code>ActivationKeyExpired</code>	The specified activation key has expired.	ActivateGateway
<code>ActivationKeyInvalid</code>	The specified activation key is invalid.	ActivateGateway
<code>ActivationKeyNotFound</code>	The specified activation key was not found.	ActivateGateway
<code>BandwidthThrottleScheduleNotFound</code>	The specified bandwidth throttle was not found.	DeleteBandwidthRateLimit
<code>CannotExportSnapshot</code>	The specified snapshot cannot be exported.	CreateCachediSCSIVolume CreateStorediSCSIVolume
<code>InitiatorNotFound</code>	The specified initiator was not found.	DeleteChapCredentials
<code>DiskAlreadyAllocated</code>	The specified disk is already allocated.	AddCache AddUploadBuffer AddWorkingStorage CreateStorediSCSIVolume
<code>DiskDoesNotExist</code>	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 S3 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
Added throughput and optimization guidance	The Performance and optimization chapter now includes recommendations and best practices for maximizing the throughput of your S3 File Gateway and optimizing your deployment for SQL Server database backup use cases. For more information, see Maximizing S3 File Gateway throughput , and Optimizing S3 File Gateway for SQL Server database backups .	June 13, 2025
Added cache reports functionality	S3 File Gateway can now generate a report of the metadata for files that are currently in the local upload cache for a specific file share. For more information, see Create a cache report for your S3 File Gateway , View and manage cache reports	March 31, 2025

[for your S3 File Gateway](#),
and [Understanding the
information provided in S3
File Gateway cache reports](#).

[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 support for Dual-layer Server-Side Encryption with AWS KMS keys \(DSSE-KMS\)](#)

You can now use dual-layer server-side encryption with AWS KMS keys to encrypt the files that S3 File Gateway uploads to Amazon S3. For more information about DSSE-KMS, see [Encrypt objects stored by File Gateway in Amazon S3](#).

September 13, 2024

[Added option to turn maintenance updates on or off](#)

Storage Gateway receives regular maintenance updates that can include operating 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](#).

June 6, 2024

[Added new SMB security level](#)

S3 File Gateway now supports an additional security level that you can use to enforce 256-bit AES encryption for SMB client connections. For more information, see [Set a security level for your gateway](#).

May 23, 2024

[Gateway appliance software version reporting and release notes for S3 File Gateway](#)

Added release notes to describe the new and updated features, improvements, and fixes that are included with each version of the Amazon S3 File Gateway appliance. You can determine a gateway's software version number from the Storage Gateway console or by using the AWS CLI. For more information, see [Release Notes — Gateway Appliance Software](#).

October 5, 2023

[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

[Increased maximum file shares per gateway](#)

S3 File Gateway now supports up to 50 file shares per gateway, increased from the previous limit of 10. This allows you to create more file shares on a single gateway, reducing the number of gateways you need to manage. For more information, see [Quotas for file shares](#).

January 18, 2023

[Added support for DOS attributes](#)

S3 File Gateway now supports DOS attributes for files stored in Amazon S3. This allows you to preserve Windows file attributes such as read-only, hidden, system, and archive when files are uploaded to Amazon S3. For more information, see [Support for file attributes in Amazon S3 File Gateway](#).

January 18, 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 schedule-based network bandwidth throttling](#)

S3 File Gateway now supports schedule-based network bandwidth throttling for data uploads to Amazon S3. This feature allows you to limit the amount of network bandwidth that your gateway uses during specific time periods, helping you manage network usage during peak business hours. For more information, see [Managing bandwidth for your S3 File Gateway](#).

January 18, 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

[Support for force-closing files on SMB file shares](#)

You can now use Local Group settings to assign Gateway Admin permissions. Gateway Admins can use the Shared Folders Microsoft Management Console snap-in to force-close files that are open and locked on SMB file shares. For more information, see [Configure Local Groups for your gateway](#).

October 12, 2021

[Audit log support for NFS file shares](#)

You can now configure NFS file shares to generate audit logs that provide details about user access to files and folders within a file share. You can use these logs to monitor user activities and take action if inappropriate activity patterns are identified. For more information, see [Understanding File Gateway audit logs](#).

October 12, 2021

[Access point alias support](#)

File Gateway file shares can now connect to Amazon S3 storage using bucket-style access point aliases. For more information, see [Create a file share](#).

October 12, 2021

[VPC endpoint and access point support](#)

File Gateway file shares can now connect to S3 buckets through access points or interface endpoints in your VPC powered by AWS PrivateLink. For more information, see [Create a file share](#).

July 7, 2021

[Opportunistic locking support](#)

File Gateway file shares can now use opportunistic locking to optimize their file buffering strategy, which improves performance in most cases, particularly with regard to Windows context menus. For more information, see [Create an SMB file share](#).

July 7, 2021

FedRAMP compliance	Storage Gateway is now FedRAMP compliant. For more information, see Compliance validation for Storage Gateway .	November 24, 2020
File upload notification for File Gateway	File Gateway now provides file upload notification, which notifies you when a file has been fully uploaded to Amazon S3 by the File Gateway. For more information, see Getting file upload notification .	November 9, 2020
Access-based enumeration for File Gateway	File Gateway now provides access-based enumeration, which filters the enumeration of files and folders on an SMB file share based on the share's ACLs. For more information, see Creating an SMB file share .	November 9, 2020
File Gateway migration	File Gateway now provides a documented process for replacing an existing File Gateway with a new File Gateway. For more information, 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 information, see [Using the AWS Storage Gateway Hardware Appliance](#).

August 12, 2020

[Support for Federal Information Processing Standard \(FIPS\) endpoints in new AWS Regions](#)

You can now activate a gateway with FIPS endpoints in the US East (Ohio), US East (N. Virginia), US West (N. California), US West (Oregon), and Canada (Central) Regions. For more information, see [AWS Storage Gateway endpoints and quotas](#) in the *AWS General Reference*.

July 31, 2020

[Support for multiple file shares attached to a single Amazon S3 bucket](#)

File Gateway now supports creating multiple file shares for a single S3 bucket and synchronizing the File Gateway's local cache with a bucket based on frequency of directory access. You can limit the number of buckets necessary to manage the file shares that you create on your File Gateway. You can define multiple S3 prefixes for an S3 bucket and map a single S3 prefix to a single gateway file share. You can also define gateway file share names to be independent of the bucket name to fit the on-premises file share naming convention. For more information, see [Creating an NFS file share](#) or [Creating an SMB file share](#).

July 7, 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 CloudWatch alarms in the Storage Gateway console](#)

You can now view CloudWatch alarms in the Storage Gateway console. For more information, see [Understanding 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

[SMB file share support for Microsoft Windows ACLs](#)

For File Gateways, you can now use Microsoft Windows access control lists (ACLs) to control access to Server Message Block (SMB) file shares. For more information, see [Using Microsoft Windows ACLs to Control Access to an SMB File Share](#).

May 8, 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-gigabit 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 Hardware Appliance](#).

September 18, 2018

[Support for Server Message Block \(SMB\) protocol](#)

File Gateways added support for the Server Message Block (SMB) protocol to file shares. For more information, see [Creating a File Share](#).

June 20, 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
Support for S3 One Zone-IA storage class	For File Gateways, you can now choose S3 One Zone-IA as the default storage class for your file shares. Using this storage class, you can store your object data in a single Availability Zone in Amazon S3. For more information, see Creating a file share .	April 4, 2018
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
Support for refresh cache notification, Requester Pays, and canned ACLs for Amazon S3 buckets	<p>With File Gateways, you can now be notified when the gateway finishes refreshing the cache for your Amazon S3 bucket. For more information, see RefreshCache.html in the <i>Storage Gateway API Reference</i>.</p> <p>For File Gateways, you can now specify that the requester or reader pays for access charges instead of the bucket owner.</p> <p>File Gateway can now give full control of written files to the owner of the S3 bucket that maps to the NFS file share.</p> <p>For more information, see Creating a file share.</p>	March 1, 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 file upload notificat	File Gateway can now send notifications when all files written to your NFS file share have been	November 21, 2017

Change	Description	Date Changed
ion and guessing of the Multipurpose Internet Mail Extension (MIME) type	<p>uploaded to Amazon S3. For more information, see NotifyWhenUploaded in the <i>Storage Gateway API Reference</i>.</p> <p>File Gateway can now guess the MIME type for uploaded objects based on file extensions. For more information, see Creating a file share.</p>	
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
<p>Updates to file share settings</p> <p>Support for cache refresh for file shares</p>	<p>File Gateways now add mount options to the file share settings. You can now set squash and read-only options for your file share. For more information, see Creating a file share.</p> <p>File Gateways now can find objects in the Amazon S3 bucket that were added or removed since the gateway last listed the bucket's contents and cached the results. For more information, see RefreshCache in the API Reference.</p>	March 28, 2017

Change	Description	Date Changed
Support for File Gateways on Amazon EC2	<p>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 S3 File Gateway. For information about how to launch a File Gateway AMI, see Deploy a default Amazon EC2 host for S3 File Gateway.</p> <p>In addition, File Gateway now supports HTTP proxy configuration. For more information, see Routing your gateway deployed on Amazon EC2 through an HTTP proxy.</p>	February 08, 2017
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

Release notes for gateway appliance software

These release notes describe the new and updated features, improvements, and fixes that are included with each version of the Amazon S3 File Gateway appliance. Each software version is identified by its release date and a unique version number.

You can determine a gateway's software version number by checking its **Details** page in the Storage Gateway console, or by calling the [DescribeGatewayInformation](#) API action using an AWS CLI command similar to the following:

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

The version number is returned in the `SoftwareVersion` field of the API response.

Note

A gateway won't report software version information under the following circumstances:

- The gateway is offline.
- The gateway is running older software that doesn't support version reporting.
- The gateway type isn't S3 File Gateway.

For more information about S3 File Gateway updates, including how to modify the default automatic maintenance and update schedule for a gateway, see [Managing Gateway Updates Using the AWS Storage Gateway Console](#).

Amazon Linux 2023 (AL2023) based gateways

The following table lists the release notes for gateways based on AL2023.

Note

Gateway versions 1.x.x can't be updated to 2.x.x.


Release Date	Software Version	Release Notes
2025-08-29	2.0.1	<ul style="list-style-type: none">Updated operating system and software elements to improve security and performance.Initial release for on-premises gateways.
2025-08-21	2.0.0	<ul style="list-style-type: none">Initial release of new operating system.Added IPv6 support.

Amazon Linux 2 (AL2) based gateways

The following table lists the release notes for gateways based on AL2.

Release Date	Software Version	Release Notes
2025-08-21	1.27.13	Maintenance Updates: <ul style="list-style-type: none">Added system statistics and metrics to provide deeper insights into gateway performance.
2025-08-18	1.27.12	Maintenance Updates: <ul style="list-style-type: none">Updated operating system and software elements to improve security and performance.
2025-08-11	1.27.11	Maintenance Updates: <ul style="list-style-type: none">Fixed an issue affecting S3 metadata updates under

Release Date	Software Version	Release Notes
		specific file operations for some gateways.
2025-07-15	1.27.10	Maintenance Updates: <ul style="list-style-type: none">• Updated operating system and software elements to improve security and performance.• Resolved issues for Upload Rename Order log for gateways.
2025-06-23	1.27.9	Maintenance Updates: <ul style="list-style-type: none">• Resolved issues for Upload Rename Order log for gateways.• Enabled Upload Rename Order log for new gateways.• Fixed issue so that the gateway now correctly handles delete operations for files that were not successfully uploaded.
2025-06-16	1.27.8	Maintenance Updates: <ul style="list-style-type: none">• Updated operating system and software elements to improve security and performance.

Release Date	Software Version	Release Notes
2025-05-26	1.27.7	<p>Maintenance Updates:</p> <ul style="list-style-type: none">Resolved issues for Rename Ordering. <div> Note This issue only affects some gateways. You will be notified if your gateway needs to be updated.</div> <ul style="list-style-type: none">Updated operating system and software elements to improve security and performance.
2025-05-15	1.27.6	<p>Maintenance Updates:</p> <ul style="list-style-type: none">Updated operating system and software elements to improve security and performance.

Release Date	Software Version	Release Notes
2025-04-28	1.27.5	Maintenance Updates: <ul style="list-style-type: none">• Fixed issues related to the upload rename order log.• Added debug tooling to help determine the causes of upload issues.• Added system statistics and metrics for deeper insight of gateway performance.• Updated operating system and software elements to improve security and performance.
2025-04-14	1.27.4	Maintenance Updates: <ul style="list-style-type: none">• Updated operating system and software elements to improve security and performance.
2025-04-01	1.27.3	Maintenance Updates: <ul style="list-style-type: none">• Update for gateway logging configuration. No customer action required.

Release Date	Software Version	Release Notes
2025-03-17	1.27.2	<p>Maintenance Updates:</p> <ul style="list-style-type: none">• Added an API action that removes cached file share data and metadata for files that have failed to upload from your gateway to Amazon S3. Your AWS account must be allowlisted to use this function. If your gateway has problems with files failing upload, AWS Support can unlock the function and provide guidance about its use.• Added capability for new gateways to log the order of object rename upload operations. This helps prevent files from failing to upload to Amazon S3 after repeated or overlapping rename operations.• Fixed an issue related to SMB unintentionally opening ports.• Updated operating system and software elements to improve security and performance.

Release Date	Software Version	Release Notes
2025-02-17	1.27.1	Maintenance Updates: <ul style="list-style-type: none">• Removed Java 11.• Added cache function for Support use.• Updated operating system and software elements to improve security and performance.
2025-01-17	1.27.0	Features: <ul style="list-style-type: none">• Cache Reports (coming soon) - Updated gateway software to support launch of an upcoming feature designed to generate reports of the file metadata currently cached by an S3 File Gateway. You will be able to use these reports to troubleshoot issues if you have files failing upload from your gateway to Amazon S3. Maintenance Updates: <ul style="list-style-type: none">• Updated operating system and software elements to improve security and performance.

Release Date	Software Version	Release Notes
2025-01-09	1.26.9	Maintenance Updates: <ul style="list-style-type: none">Updated operating system and software elements to improve security and performance.
2024-12-18	1.26.8	Maintenance Updates: <ul style="list-style-type: none">Updated operating system and software elements to improve security and performance.
2024-11-18	1.26.7	Maintenance Updates: <ul style="list-style-type: none">Updated operating system and software elements to improve security and performance.
2024-10-17	1.26.6	Maintenance Updates: <ul style="list-style-type: none">Updated operating system and software elements to improve security and performance.
2024-09-30	1.26.5	Maintenance Updates: <ul style="list-style-type: none">Fixed an issue with on-premises gateways not allowing support channelsUpdated operating system and software elements to improve security and performance.

Release Date	Software Version	Release Notes
2024-09-16	1.26.3	Maintenance Updates: <ul style="list-style-type: none">Updated operating system and software elements to improve security and performance.
2024-08-21	1.26.1	Maintenance Updates: <ul style="list-style-type: none">Fixed an issue related to logging.
2024-08-19	1.26.0	Maintenance Updates: <ul style="list-style-type: none">Updated operating system and software elements to improve security and performance.
2024-07-16	1.25.2	Maintenance Updates: <ul style="list-style-type: none">Updated operating system and software elements to improve security and performance.
2024-06-17	1.25.1	Maintenance Updates: <ul style="list-style-type: none">Fixed an issue with upgrades when using a proxy and DNS is disabled.Updated operating system and software elements to improve security and performance.

Release Date	Software Version	Release Notes
2024-05-15	1.25.0	<p>Features:</p> <ul style="list-style-type: none">• Added ability to set AES-128 or AES-256 encryption minimum. This is a gateway change only and will be available in the Storage Gateway console in an upcoming release.• Increased rotation of system logs when disk space is low. Previously, log writing that filled the root disk would cause the gateway to stop. Now, as space decreases, the gateway will make more room for newer logs by eliminating older logs.• Added S3 path in health notifications for file upload errors. Previously, health notifications only showed the path to the file on the gateway. Notifications now show the path to help users locate the file in S3.• Service now ignores backend blockers during forced file share deletes. Previously, forced deletes would stop without explanation when encountering blockers.

Release Date	Software Version	Release Notes
		<p>Forced deletes now continue uninterrupted in these scenarios.</p> <p>Maintenance Updates:</p> <ul style="list-style-type: none">• Updated the NFS stack.• Upgraded Java 17 JRE.• Updated operating system and software elements to improve security and performance.
2024-04-15	1.24.5	<p>Maintenance Updates:</p> <ul style="list-style-type: none">• Updated operating system and software elements to improve security and performance.
2024-04-01	1.24.4	<p>Maintenance Updates:</p> <ul style="list-style-type: none">• Addressed missing Network Time Protocol (NTP) component.

Release Date	Software Version	Release Notes
2024-03-18	1.24.3	Maintenance Updates: <ul style="list-style-type: none">• Fixed an issue related to case-sensitive lookup performance.• Fixed an issue that caused processes to crash.• Updated operating system and software elements to improve security and performance.
2024-01-12	1.24.2	Maintenance Updates: <ul style="list-style-type: none">• Addressed SMB logging issue.
2023-12-27	1.24.1	Maintenance Updates: <ul style="list-style-type: none">• Addressed SMB stability issue.

Release Date	Software Version	Release Notes
2023-12-01	1.24.0	<p>Features:</p> <ul style="list-style-type: none">• Updated SMB stack.<ul style="list-style-type: none">• Added support for AES-256 encryption, plus more secure variants of AES-128 encryption and signing when using an SMB 3.1.1 client which requests it.• SMBv1 (LANMAN/CIFS) server-side copy and server-side wildcard expansion functionality have been removed. (SMBv2 and SMBv3 are unaffected.) This may negatively impact performance of certain SMBv1 workloads. If you use SMBv1, you are encouraged to migrate to SMBv2 or SMBv3. <p>Maintenance Updates:</p> <ul style="list-style-type: none">• Updated operating system and software elements to improve security and performance.

Release Date	Software Version	Release Notes
2023-10-24	1.23.2	Maintenance Updates: <ul style="list-style-type: none">Fixed an issue related to Support Channel not connecting properly for certain users.
2023-08-14	1.23.1	Maintenance Updates: <ul style="list-style-type: none">Updated NTP server to use sync server for new gateways.
2023-06-12	1.23.0	Features: <ul style="list-style-type: none">Increased upload threads for some AWS accounts. Maintenance Updates: <ul style="list-style-type: none">Fixed an access violation issue on large copies.Fixed an NFS issue.Removed Java 8.Updated operating system and software elements to improve security and performance.
2023-04-19	1.22.1	Maintenance Updates: <ul style="list-style-type: none">Fixed an issue related to renaming folders and files.

Release Date	Software Version	Release Notes
2023-01-18	1.22.0	Features: <ul style="list-style-type: none">• Added support for DOS attributes.• Increased the number of supported file shares per gateway from 10 to 50.• Implemented a clock skew detection mechanism to determine when a gateway and a service are out of sync. Maintenance Updates: <ul style="list-style-type: none">• Updated SMB stack.
2022-07-06	1.21.2	Maintenance Updates: <ul style="list-style-type: none">• Updated operating system and software elements to improve security and performance.
2022-02-16	1.21.1	Features: <ul style="list-style-type: none">• Added new metrics for rename and deletion in cache. Maintenance Updates: <ul style="list-style-type: none">• Fixed miscellaneous issues.

Release Date	Software Version	Release Notes
2022-01-18	1.21.0	Features: <ul style="list-style-type: none">Added new CloudWatch metrics.Added bandwidth throttling for data uploads.
2021-12-12	1.20.0	<i>URGENT UPDATE:</i> <ul style="list-style-type: none">Addressed Log4j vulnerability.