

### Amazon S3 File Gateway User Guide

# **Amazon Storage Gateway**



## Amazon Storage Gateway: Amazon S3 File Gateway User Guide

## **Table of Contents**

| ••• |  | , xi |
|-----|--|------|
| W   | hat is Amazon S3 File Gateway                                | 1    |
|     | How S3 File Gateway works                                    | 2    |
| G   | etting Started   | 5    |
|     | Sign up for Amazon Web Services                              | . 5  |
|     | Create an IAM user with administrator privileges             | 5    |
|     | Secure IAM users   | 5    |
|     | Requirements   | 6    |
|     | Prerequisites  | 7    |
|     | Hardware and storage requirements                            | . 7  |
|     | Network and firewall requirements                            | . 9  |
|     | Supported hypervisors and host requirements                  | 24   |
|     | Supported NFS clients for File Gateway                       | 25   |
|     | Supported SMB clients for File Gateway                       | 26   |
|     | Supported file system operations                             | 26   |
|     | Accessing Amazon Storage Gateway                             | 26   |
|     | Supported Amazon Regions                                     | 27   |
| Us  | sing the hardware appliance                                  | 28   |
|     | Ordering information   | 28   |
|     | Supported Amazon Regions                                     | 29   |
|     | Setting up your hardware appliance                           | 29   |
|     | Rack-mounting and connecting the hardware appliance to power | 30   |
|     | Hardware appliance dimensions                                | 30   |
|     | Configuring network parameters                               | 35   |
|     | Activating your hardware appliance                           | 38   |
|     | Creating a gateway   | 39   |
|     | Configuring an IP address for the gateway                    | 40   |
|     | Configuring your gateway                                     | 42   |
|     | Removing a gateway   | 42   |
|     | Deleting your hardware appliance                             | 43   |
| Cr  | eating Your Gateway  | 44   |
|     | Overview - Gateway Activation                                | 44   |
|     | Set up gateway   | 44   |
|     | Connect to Amazon  | 44   |

|    | Review and activate   | 45   |
|----|---|------|
|    | Overview - Gateway Configuration                            | 45   |
|    | Overview - Storage Resources                                | 45   |
|    | Create an S3 File Gateway                                   | 45   |
|    | Set up an Amazon S3 File Gateway                            | 46   |
|    | Connect your Amazon S3 File Gateway to Amazon               | 47   |
|    | Review settings and activate your Amazon S3 File Gateway    | 48   |
|    | Configure your Amazon S3 File Gateway                       | 48   |
|    | Activating a gateway in a VPC                               | 50   |
|    | Creating a VPC endpoint for Storage Gateway                 | 51   |
| Cr | eate a file share   | 53   |
|    | Avoid unanticipated costs                                   | 54   |
|    | Encrypt objects stored by File Gateway                      | 55   |
|    | Create NFS file share with default settings                 | . 55 |
|    | Create an NFS file share                                    | 55   |
|    | NFS file share default settings                             | 56   |
|    | Create an NFS file share with custom configuration          | 59   |
|    | Create SMB file share with default settings                 | 65   |
|    | Create an SMB file share                                    | 66   |
|    | SMB file share default settings                             | . 56 |
|    | Create an SMB file share with a custom configuration        | 70   |
|    | Creating an SMB file share                                  | 71   |
| Μ  | ount and use your file share                                | . 80 |
|    | Mount your NFS file share on your client                    | . 80 |
|    | Mount your SMB file share on your client                    | 82   |
|    | Using file shares on buckets with pre-existing objects      | 87   |
|    | Test your S3 File Gateway                                   | . 87 |
|    | Cleaning up resources you don't need                        | 88   |
| M  | anaging your Amazon S3 File Gateway                         | 90   |
|    | Editing Gateway Information                                 | 90   |
|    | Granting access and permissions for file shares and buckets | 91   |
|    | Granting access to an S3 bucket                             | 91   |
|    | Cross-service confused deputy prevention                    | 94   |
|    | Using a file share for cross-account access                 | 96   |
|    | Deleting a file share                                       | 97   |
|    | Editing SMB settings for a gateway                          | 98   |
|    |   |      |

| Setting a security level for your gateway                                       | 99  |
|---|-----|
| Using Active Directory to authenticate users                                    | 100 |
| Providing guest access to your file share                                       | 102 |
| Configure Local Groups for your gateway   | 103 |
| Setting file share visibility   | 104 |
| Editing settings for your NFS file share  | 104 |
| Editing the metadata defaults for your NFS file share                           | 105 |
| Editing access settings for your NFS file share                                 | 107 |
| Editing settings for your SMB file share  | 107 |
| Editing access settings for your SMB file share                                 | 109 |
| Refresh Amazon S3 bucket objects  | 110 |
| Configure an automated cache refresh schedule using the Storage Gateway console | 110 |
| Configure an automated cache refresh schedule using Amazon Lambda with an Amazo | on  |
| CloudWatch rule   | 111 |
| Perform a manual cache refresh using the Storage Gateway console                | 114 |
| Perform a manual cache refresh using the Storage Gateway APIAPI                 | 114 |
| Using S3 Object Lock with an Amazon S3 File Gateway                             | 115 |
| Understanding file share status   | 115 |
| Understanding gateway status  | 116 |
| Monitoring Storage Gateway  | 118 |
| Understanding CloudWatch alarms   | 118 |
| Create recommended CloudWatch alarms  | 120 |
| Creating a custom CloudWatch alarm  | 121 |
| Monitoring your S3 File Gateway   | 123 |
| Getting S3 File Gateway health logs   | 123 |
| Using Amazon CloudWatch metrics   | 126 |
| Getting notified about file operations  | 127 |
| Understanding gateway metrics   | 135 |
| Understanding file share metrics  | 141 |
| Understanding S3 File Gateway audit logs  | 144 |
| Maintaining your gateway  | 150 |
| Shutting down your gateway VM   | 150 |
| Managing local disks  | 150 |
| Deciding the amount of local disk storage                                       | 151 |
| Sizing cache storage  | 151 |
| Add cache storage   | 152 |

| Managing Bandwidth   |    | Using ephemeral storage with EC2 gateways                               | 153 |
|--|----|---|-----|
| Using the Amazon SDK for Java  |    | Managing Bandwidth  | 154 |
| Using the Amazon SDK for .NET  |    | Edit bandwidth-rate-limit schedule                                      | 155 |
| Using the Amazon Tools for Windows PowerShell  |    | Using the Amazon SDK for Java   | 156 |
| Managing Gateway Updates   |    | Using the Amazon SDK for .NET   | 159 |
| Performing maintenance tasks using the local console   |    | Using the Amazon Tools for Windows PowerShell                           | 161 |
| Accessing the Gateway Local Console  |    | Managing Gateway Updates  | 162 |
| Performing tasks on the virtual machine local console  |    | Performing maintenance tasks using the local console                    | 164 |
| Performing tasks on the EC2 local console  |    | Accessing the Gateway Local Console                                     | 164 |
| Configuring Network Adapters for Your Gateway  |    | Performing tasks on the virtual machine local console                   | 170 |
| Deleting Your Gateway and Removing Resources   |    | Performing tasks on the EC2 local console                               | 185 |
| Deleting Your Gateway by Using the Storage Gateway Console   |    | Configuring Network Adapters for Your Gateway                           | 192 |
| Removing Resources from a Gateway Deployed On-Premises   |    |   |     |
| Removing Resources from a Gateway Deployed On-Premises   |    | Deleting Your Gateway by Using the Storage Gateway Console              | 197 |
| Removing Resources from a Gateway Deployed on an Amazon EC2 Instance 20  Replacing your existing File Gateway with a new instance 20  Method 1: Migrate cache disk and Gateway ID to replacement instance 20  Method 2: Replacement instance with empty cache disk and new Gateway ID 20  Performance and optimization 20  Basic performance guidance for S3 File Gateway 20  S3 File Gateway performance on Linux clients 20  File Gateway performance on Windows clients 21  Performance guidance for gateways with multiple file shares 21  Optimizing Gateway Performance 21  Add Resources to Your Gateway 21  Add Resources to Your Application Environment 21  Using VMware High Availability with Storage Gateway 21  Configure Your vSphere VMware HA Cluster 21  Set Up Your Gateway Type 21  Deploy the Gateway Type 21  Optional) Add Override Options for Other VMs on Your Cluster 21  Activate Your Gateway 22  Test Your VMware High Availability Configuration 22  Security 22  Security 22  Security 22  Security 20  Security 20  Activate Your VMware High Availability Configuration 22  Security 22  Security 22  Security 22  Security 22  Security 22  Security 24  Security 24  Security 25  Security 26  Security 26  Security 26  Security 27  Security 28  Security 28  Security 29  Security 20  S |    |   |     |
| Replacing your existing File Gateway with a new instance20Method 1: Migrate cache disk and Gateway ID to replacement instance20Method 2: Replacement instance with empty cache disk and new Gateway ID20Performance and optimization20Basic performance guidance for S3 File Gateway20S3 File Gateway performance on Linux clients20File Gateway performance on Windows clients21Performance guidance for gateways with multiple file shares21Optimizing Gateway Performance21Add Resources to Your Gateway21Add Resources to Your Application Environment21Using VMware High Availability with Storage Gateway21Configure Your vSphere VMware HA Cluster21Set Up Your Gateway Type21Deploy the Gateway21(Optional) Add Override Options for Other VMs on Your Cluster21Activate Your Gateway22Test Your VMware High Availability Configuration22Security22  |    |   |     |
| Method 2: Replacement instance with empty cache disk and new Gateway ID20Performance and optimization20Basic performance guidance for S3 File Gateway20S3 File Gateway performance on Linux clients20File Gateway performance on Windows clients21Performance guidance for gateways with multiple file shares21Optimizing Gateway Performance21Add Resources to Your Gateway21Add Resources to Your Application Environment21Using VMware High Availability with Storage Gateway21Configure Your vSphere VMware HA Cluster21Set Up Your Gateway Type21Deploy the Gateway21(Optional) Add Override Options for Other VMs on Your Cluster21Activate Your Gateway22Test Your VMware High Availability Configuration22Security22   | Re |   |     |
| Performance and optimization   |    | Method 1: Migrate cache disk and Gateway ID to replacement instance     | 201 |
| Basic performance guidance for S3 File Gateway   |    | Method 2: Replacement instance with empty cache disk and new Gateway ID | 204 |
| S3 File Gateway performance on Linux clients   | Pe | rformance and optimization  | 207 |
| File Gateway performance on Windows clients  |    | Basic performance guidance for S3 File Gateway                          | 207 |
| Performance guidance for gateways with multiple file shares  |    | S3 File Gateway performance on Linux clients                            | 208 |
| Optimizing Gateway Performance   |    | File Gateway performance on Windows clients                             | 210 |
| Add Resources to Your Gateway  |    | Performance guidance for gateways with multiple file shares             | 211 |
| Add Resources to Your Application Environment  |    | Optimizing Gateway Performance  | 213 |
| Using VMware High Availability with Storage Gateway  |    | Add Resources to Your Gateway   | 214 |
| Configure Your vSphere VMware HA Cluster   |    | Add Resources to Your Application Environment                           | 216 |
| Set Up Your Gateway Type   |    | Using VMware High Availability with Storage Gateway                     | 216 |
| Deploy the Gateway   |    | Configure Your vSphere VMware HA Cluster                                | 217 |
| (Optional) Add Override Options for Other VMs on Your Cluster  |    | Set Up Your Gateway Type  | 218 |
| Activate Your Gateway  |    | Deploy the Gateway  | 218 |
| Test Your VMware High Availability Configuration   |    | (Optional) Add Override Options for Other VMs on Your Cluster           | 219 |
| Test Your VMware High Availability Configuration   |    | Activate Your Gateway   | 220 |
| Security 22  |    |   |     |
|  | Se | curity  | 221 |
| ==:  |    |   |     |

|    | Data encryption   | 223   |
|----|---|-------|
|    | Identity and access management  | . 223 |
|    | Audience  | 224   |
|    | Authenticating with identities  | 224   |
|    | Managing access using policies  | 227   |
|    | How Amazon Storage Gateway works with IAM                                       | . 230 |
|    | Identity-based policy examples  | 237   |
|    | Troubleshooting   | 240   |
|    | Using Tags to Control Access to Resources                                       | . 242 |
|    | Using ACLs for SMB File Share Access  | 244   |
|    | Logging and monitoring  | 248   |
|    | Storage Gateway information in CloudTrail                                       | 248   |
|    | Understanding Storage Gateway log file entries                                  | 249   |
|    | Compliance validation   | 251   |
|    | Resilience  | 252   |
|    | Infrastructure security   | . 252 |
|    | Amazon Security Best Practices  | . 253 |
| Tr | oubleshooting   | 254   |
|    | Troubleshooting: on-premises gateway issues                                     | 254   |
|    | Turning on Amazon Web Services Support access to help troubleshoot your gateway | 258   |
|    | Troubleshooting: Microsoft Hyper-V setup issues                                 | 259   |
|    | Troubleshooting: Amazon EC2 gateway issues                                      | . 263 |
|    | Gateway activation hasn't occurred after a few moments                          | 264   |
|    | Can't find the EC2 gateway instance in the instance list                        | 264   |
|    | Connect to your Amazon EC2 gateway using the serial console                     | 265   |
|    | Turning on Amazon Web Services Support access to help troubleshoot the gateway  | 265   |
|    | Troubleshooting: hardware appliance issues                                      | 267   |
|    | How to determine service IP address   | . 267 |
|    | How to perform a factory reset  | 267   |
|    | How to perform a remote restart   | . 267 |
|    | How to obtain Dell iDRAC support  | . 268 |
|    | How to find the hardware appliance serial number                                | 268   |
|    | How to get hardware appliance support   | . 268 |
|    | Troubleshooting: File Gateway issues  | . 269 |
|    | Error: 1344 (0x00000540)  | 270   |
|    | Error: GatewayClockOutOfSync  | 270   |
|    |   |       |

| Error: InaccessibleStorageClass   | 271       |
|---|-----------|
| Error: InvalidObjectState   | 271       |
| Error: ObjectMissing  | 272       |
| Error: RoleTrustRelationshipInvalid   | 272       |
| Error: S3AccessDenied   | 272       |
| Error: DroppedNotifications   | 273       |
| Notification: HardReboot  | 274       |
| Notification: Reboot  | 274       |
| Troubleshooting with CloudWatch metrics                                       | 274       |
| Troubleshooting: file share issues  | 277       |
| File share is stuck in CREATING status  | 277       |
| Can't create a file share   | 278       |
| SMB file shares don't allow multiple different access methods                 | 278       |
| Multiple file shares can't write to the mapped S3 bucket                      | 279       |
| Notification for deleted log group when using audit logs                      | 279       |
| Can't upload files into S3 bucket   | 279       |
| Can't change default encryption to SSE-KMS                                    | 279       |
| Changes made directly in an S3 bucket with object versioning turned on may af | fect what |
| you see in your file share  | 280       |
| When writing to an S3 bucket with versioning turned on, the Amazon S3 File Ga |           |
| create multiple versions of Amazon S3 objects                                 | 281       |
| Changes to an S3 bucket are not reflected in Storage Gateway                  | 282       |
| ACL permissions aren't working as expected                                    | 283       |
| Gateway performance declined after a recursive operation                      |           |
| High Availability Health Notifications  | 283       |
| Troubleshooting: high availability issues                                     | 283       |
| Health notifications  | 284       |
| Metrics   | 285       |
| Best Practices  | 286       |
| Best practices: recovering data   | 286       |
| Recovering from an unexpected VM shutdown                                     | 287       |
| Recovering data from a malfunctioning cache disk                              | 287       |
| Recovering data from an inaccessible data center                              | 287       |
| Best practices: multipart uploads   |           |
| Unzip compressed files locally  |           |
| Copying data from Windows Server  | 289       |

|    | Best practices: Proper cache disk size                  | 289 |
|----|---|-----|
|    | Working with multiple file shares and Amazon S3 buckets | 289 |
|    | Allowing specific NFS clients to mount your file share  | 291 |
| Ac | Iditional Resources                                     | 292 |
|    | Host setup  | 292 |
|    | Configuring VMware for Storage Gateway                  | 293 |
|    | Synchronizing Your Gateway VM Time                      | 298 |
|    | Deploy an Amazon EC2 host for File Gateway              | 300 |
|    | Deploy an Amazon EC2 with Default Settings              | 304 |
|    | Modify Amazon EC2 instance metadata options             | 306 |
|    | Getting Activation Key                                  | 307 |
|    | Linux (curl)  | 308 |
|    | Linux (bash/zsh)  | 309 |
|    | Microsoft Windows PowerShell                            | 309 |
|    | Using your local console                                | 310 |
|    | DOS or Windows file attribute support                   | 310 |
|    | Using Amazon Direct Connect with Storage Gateway        | 311 |
|    | Port Requirements                                       | 312 |
|    | Active Directory permissions                            | 322 |
|    | Connecting to Your Gateway                              | 323 |
|    | Getting an IP Address from an Amazon EC2 Host           | 324 |
|    | Understanding Resources and Resource IDs                | 325 |
|    | Working with Resource IDs                               | 325 |
|    | Tagging Your Resources                                  | 326 |
|    | Working with tags                                       | 327 |
|    | See also  | 328 |
|    | Open-source components                                  | 328 |
|    | Open-source components for Storage Gateway              | 328 |
|    | Open-source components for Amazon S3 File Gateway       | 329 |
|    | Quotas  | 329 |
|    | Quotas for file shares                                  | 329 |
|    | Recommended local disk sizes for your gateway           | 330 |
|    | Using storage classes                                   | 331 |
|    | Using storage classes with a File Gateway               | 331 |
|    | Using the GLACIER storage class with File Gateway       | 336 |
|    | Using Kubernetes CSI drivers                            | 336 |

| SMB CSI drivers               | 337 |
|-------------------------------|-----|
| NFS CSI drivers               | 341 |
| Terraform module              | 345 |
| API Reference                 | 347 |
| Required Request Headers      | 347 |
| Signing Requests              | 349 |
| Example Signature Calculation | 350 |
| Error Responses               | 352 |
| Exceptions                    | 353 |
| Operation Error Codes         | 354 |
| Error Responses               | 375 |
| Actions                       | 377 |
| Document history              | 378 |
| Earlier updates               |     |
| Release Notes                 | 394 |
|                               |     |

Amazon FSx File Gateway documentation has been moved to What is Amazon FSx File Gateway?

Volume Gateway documentation has been moved to What is Volume Gateway?

Tape Gateway documentation has been moved to What is Tape Gateway?

## 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 Amazon, 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 Amazon 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 Amazon, buffers applications from network congestion, optimizes and streams data in parallel, and manages bandwidth consumption.

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

- Common access management using Amazon Identity and Access Management (IAM)
- Encryption using Amazon Key Management Service (Amazon KMS)
- Monitoring using Amazon CloudWatch (CloudWatch)
- Audit using Amazon CloudTrail (CloudTrail)
- Operations using the Amazon Web Services Management Console and Amazon Command Line Interface (Amazon 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 Amazon Web Services Management Console. If you want to perform these operations programmatically, see the *Amazon 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 Amazon Web Services 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.

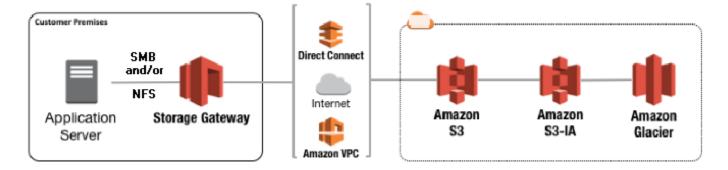
When sending HTTPS data upload requests to the Amazon S3, File Gateway populates the Content-MD5 header with the MD5 checksum of the data being uploaded. The use of this header causes Amazon S3 to return a failure in case of mismatch between the MD5 checksum computed by Amazon S3 and the value received 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 Amazon using multipart parallel uploads or byte-range downloads, to better use the available bandwidth. Local cache is maintained

How S3 File Gateway works API Version 2013-06-30 2

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 Amazon. 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 Amazon 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<br>objects and creates new S3<br>objects for each folder and<br>files in the folder structure | Early deletion fees and retrieval fees may apply |
| Change file/folder permissions | Replaces existing S3 object<br>and creates a new S3 object<br>for each file or folder                                  | Early deletion fees and retrieval fees may apply |

How S3 File Gateway works API Version 2013-06-30 3

| File operation               | S3 object impact  | Storage class implication                        |
|------------------------------|---|--|
| Change file/folder ownership | Replaces existing S3 object<br>and creates a new S3 object<br>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.

How S3 File Gateway works API Version 2013-06-30 4

## **Getting Started**

This section provides instructions for getting started with Amazon S3 File Gateway. To get started, you first sign up for Amazon. If you are a first-time user, we recommend that you read the <u>Regions</u> and <u>Requirements</u> sections.

#### **Topics**

- Sign up for Amazon Web Services
- Create an IAM user with administrator privileges
- File Gateway setup requirements
- Accessing Amazon Storage Gateway
- Supported Amazon Regions

### Sign up for Amazon Web Services

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

#### To sign up for Amazon Web Services

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

### Create an IAM user with administrator privileges

After you create your Amazon account, use the following steps to create an Amazon Identity and Access Management (IAM) user for yourself. Then you add that user to a group that has administrative permissions.

### **Secure IAM users**

After you sign up for an Amazon Web Services account, safeguard your administrative user by turning on multi-factor authentication (MFA). For instructions, see <a href="Enable a virtual MFA device for an IAM user (console)">Enable a virtual MFA device for an IAM user (console)</a> in the IAM User Guide.

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

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

- Creating an IAM user in your Amazon Web Services account
- Access management for Amazon resources
- Example IAM identity-based policies

#### Marning

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.

## File Gateway setup requirements

Unless otherwise noted, the following requirements are common to all File Gateway types in Amazon 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 clients for File Gateway
- Supported SMB clients for File Gateway
- Supported file system operations for File Gateway

Requirements API Version 2013-06-30 6

### **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 <u>Active Directory service account permission</u> requirements.
- Ensure that there is sufficient network bandwidth between the gateway and Amazon. A minimum of 100 Mbps is required to successfully download, activate, and update the gateway.
- Configure your private networking, VPN, or Amazon Direct Connect between your Amazon
  Virtual Private Cloud (Amazon VPC) and the on-premises environment where you are deploying
  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 <u>Basic performance</u> 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 Optimizing Gateway Performance. For information about how your hardware affects the performance of the gateway VM, see Quotas for file shares.

Prerequisites API Version 2013-06-30 7

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

#### **Recommended for File Gateway types**

• General-purpose instance family – **m4**, **m5**, **or m6** instance type.



#### Note

We don't recommend using the **m4.16xlarge** instance type.

- Compute-optimized instance family c4, c5, or c6 instance types. Choose the 2xlarge instance size or higher to meet the required RAM requirements.
- Memory-optimized instance family r3, r5, or r6 instance types.
- Storage-optimized instance family i3 or i4 instance types.

#### 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      | Cache     | Cache     |
|--------------|-----------|-----------|
| type         | (minimum) | (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.



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 Amazon IP address ranges. In these cases, your gateway might experience service connectivity issues when the Amazon IP range values changes. The Amazon IP address range values that you need to use are in the Amazon service subset for the Amazon Region that you activate your gateway in. For the current IP range values, see Amazon IP address ranges in the Amazon Web Services General Reference.

#### **Topics**

- · Port requirements
- Networking and firewall requirements for the Storage Gateway Hardware Appliance
- Allowing Amazon Storage Gateway access through firewalls and routers
- Configuring security groups for your Amazon EC2 gateway instance

#### Port requirements

Storage Gateway requires certain ports to be allowed for its operation. The following illustrations show the required ports that you must allow for each type of gateway. Some ports are required by all gateway types, and others are required by specific gateway types. For more information about port requirements, see Port Requirements.

#### Common ports for all gateway types

The following ports are common to all gateway types and are required by all gateway types.

| Protocol | Port        | Direction | Source             | Destination | Usage  |
|----------|-------------|-----------|--------------------|-------------|--|
| TCP      | 443 (HTTPS) | Outbound  | Storage<br>Gateway | Amazon      | For<br>communica<br>tion from<br>Storage<br>Gateway to<br>the Amazon |

| Protocol | Port | Direction | Source | Destination | Usage  |
|----------|------|-----------|--------|-------------|--|
|          |      |           |        |             | service endpoint. For information about service endpoints, see Allowing Amazon Storage Gateway access through firewalls and routers. |

| Protocol | Port      | Direction | Source   | Destination     | Usage   |
|----------|-----------|-----------|--|-----------------|---|
| TCP      | 80 (HTTP) | Inbound   | The host from which you connect to the Amazon Web Services Management Console. | Storage Gateway | By local systems to obtain the Storage Gateway activation key. Port 80 is only used during activation of the Storage Gateway appliance.  Storage Gateway does not require port 80 to be publicly accessible. The required level of access to port 80 depends on your network configura tion. If you activate your gateway from the Storage Gateway console, the |

| Protocol | Port     | Direction | Source             | Destination | Usage  |
|----------|----------|-----------|--------------------|-------------|--|
|          |          |           |                    |             | host from which you connect to the console must have access to your gateway's port 80. |
| TCP/UDP  | 53 (DNS) | Outbound  | Storage<br>Gateway | DNS server  | For communica tion between Storage Gateway and the DNS server.                         |

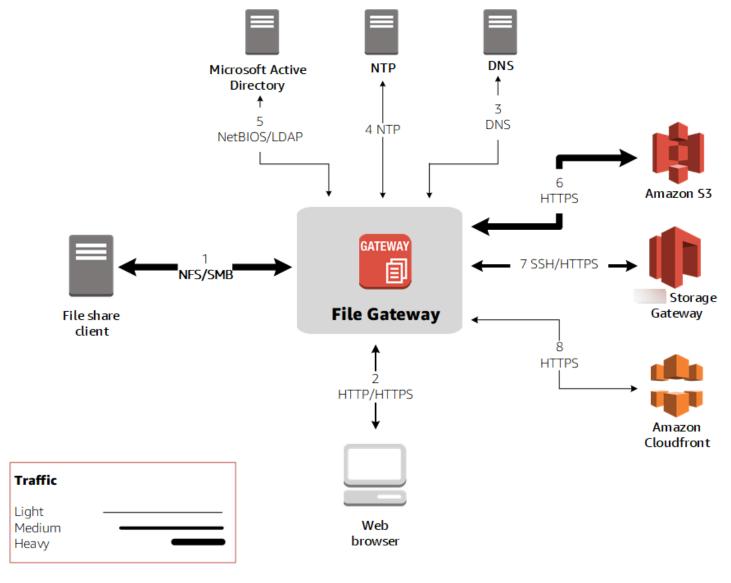
| Protocol | Port                 | Direction | Source             | Destination                       | Usage  |
|----------|----------------------|-----------|--------------------|-----------------------------------|--|
| TCP      | 22 (Support channel) | Outbound  | Storage<br>Gateway | Amazon<br>Web Services<br>Support | Allows Amazon Web Services Support to access your gateway to help you with troublesh ooting gateway issues. You don't need this port open for the normal operation of your gateway, but it is required for troublesh ooting. |

| Protocol | Port      | Direction | Source     | Destination | Usage  |
|----------|-----------|-----------|------------|-------------|--|
| UDP      | 123 (NTP) | Outbound  | NTP client | NTP server  | Used by local systems to synchroni ze VM time to the host time.  (i) Note  Not required for gateways hosted on Amazon EC2. |
| TCP      | 1026      | Outbound  |            |             | Used for control traffic.  |
| TCP      | 1027      | Outbound  |            |             | Used only during activation and can then be closed.  |
| TCP      | 1028      | Outbound  |            |             | Used for control traffic.  |

| Protocol | Port | Direction | Source | Destination | Usage   |
|----------|------|-----------|--------|-------------|---|
| TCP      | 1031 | Outbound  |        |             | Used only for software updates for file gateways. |

#### **Ports for File Gateways**

The following illustration shows the ports to open for an S3 File Gateway.





For specific port requirements, see Port Requirements.

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 Amazon Directory Service to create an Amazon Managed Microsoft AD in the Amazon Web Services Cloud. For most Amazon 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 Amazon Directory Service Administration Guide.

In addition to the common ports, Amazon S3 File Gateway requires the following ports.

| Protocol | Port        | Direction | Source       | Destination        | Usage   |
|----------|-------------|-----------|--------------|--------------------|---|
| TCP/UDP  | 2049 (NFS)  | Inbound   | NFS clients  | Storage<br>Gateway | For local<br>systems to<br>connect to<br>NFS shares<br>that your<br>gateway<br>exposes. |
| TCP/UDP  | 111 (NFSv3) | Inbound   | NFSv3 client | Storage<br>Gateway | For local systems to connect to the port mapper that your gateway exposes.              |

| Protocol        | Port             | Direction            | Source             | Destination                      | Usage   |
|-----------------|------------------|----------------------|--------------------|----------------------------------|---|
|                 |                  |                      |                    |                                  | (i) Note This port is needed only for NFSv3.  |
| TCP/UDP         | 20048<br>(NFSv3) | Inbound              | NFSv3 client       | Storage<br>Gateway               | For local systems to connect to mounts that your gateway exposes.   (i) Note  This port is needed only for NFSv3. |
| TCP/UDP<br>LDAP | 389              | Inbound and outbound | Storage<br>Gateway | Microsoft<br>Active<br>Directory | Directory System Agent (DSA); client connection   |

| Protocol   | Port | Direction            | Source             | Destination                      | Usage |
|--|------|----------------------|--------------------|----------------------------------|-------|
| TCP/UDP<br>Kerberos  | 88   | Inbound and outbound | Storage<br>Gateway | Microsoft<br>Active<br>Directory |       |
| TCP Distributed Computing Environme nt/End Point Mapper (DCE/EMAP) | 135  | Inbound and outbound | Storage<br>Gateway | Microsoft<br>Active<br>Directory |       |

# Networking and firewall requirements for the Storage Gateway Hardware Appliance

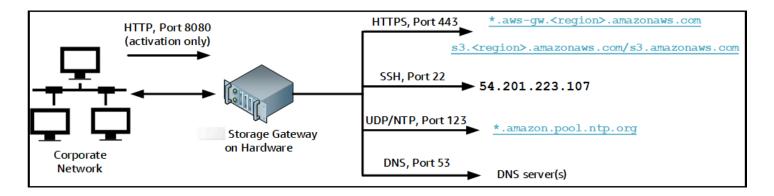
Each Storage Gateway Hardware Appliance requires the following network services:

- **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<br>appliance | 54.201.22<br>3.107      | Support channel                  |
| DNS      | 53   | Outbound  | Hardware appliance    | DNS servers             | Name<br>resolutio<br>n           |
| UDP/NTP  | 123  | Outbound  | Hardware<br>appliance | *.amazon. pool.ntp. org | Time<br>synchroni<br>zation      |
| HTTPS    | 443  | Outbound  | Hardware appliance    | *.amazona<br>ws.com     | Data<br>transfer                 |
| НТТР     | 8080 | Inbound   | Amazon                | Hardware appliance      | Activatio<br>n (only<br>briefly) |

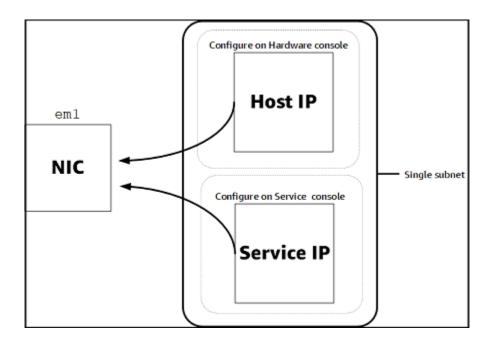
To perform as designed, a hardware appliance requires network and firewall settings as follows:

- 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 network parameters.

For an illustration showing the back of the server with its ports, see Rack-mounting your hardware appliance and connecting power.

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 Amazon Storage Gateway Hardware Appliance.

### Allowing Amazon Storage Gateway access through firewalls and routers

Your gateway requires access to the following service endpoints to communicate with Amazon. 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 Amazon.



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

#### Important

Depending on your gateway's Amazon Region, replace region in the service endpoint with the correct Region string.

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

```
s3.amazonaws.com.cn:443
```

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

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

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

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

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

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

The Amazon S3 service endpoint, shown following, is used by File Gateways only. A File Gateway requires this endpoint to access the Amazon S3 bucket that a file share maps to.

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

The following example is an S3 service endpoint in the China (Beijing) Region (cn-north-1).

```
s3.cn-north-1.amazonaws.com.cn
```



#### Note

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

A Storage Gateway VM is configured to use the following NTP servers.

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

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

### Configuring security groups for your Amazon EC2 gateway instance

In Amazon 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 Amazon Web Services Support for troubleshooting purposes. For more information, see <u>You want Amazon Web Services Support to help</u> 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 Amazon as an Amazon EC2 instance.

Storage Gateway supports the following hypervisor versions and hosts:

- VMware ESXi Hypervisor (version 7.0 or 8.0) A free version of VMware is available on the
   VMware website. 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 <u>Microsoft Download Center</u>. 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
  Deploying an Amazon EC2 instance to host an 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.



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 clients for File Gateway**

File Gateway supports the following Network File System (NFS) clients:

- Amazon Linux
- Mac OS X



#### Note

We recommend setting the rsize and wsize mount options to 1MB (or to maximum supported on client side) to improve performance when mounting NFS file shares on Mac OS X.

- RHEL 7, RHEL 8.6
- SUSE Linux Enterprise Server 11 and SUSE Linux Enterprise Server 12
- Ubuntu 14.04
- Microsoft Windows 10 Enterprise, Windows Server 2012, and Windows Server 2016. Native clients only support NFS version 3.
- Windows 7 Enterprise and Windows Server 2008.

Native clients only support NFS v3. The maximum supported NFS I/O size is 32 KB, so you might experience degraded performance on these versions of Windows.



#### Note

You can now use SMB file shares when access is required through Windows (SMB) clients instead of using Windows NFS clients.

# **Supported SMB clients for File Gateway**

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

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



## Note

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

## Supported file system operations for File Gateway

Your NFS or SMB client can write, read, delete, and truncate files. When clients send writes to Amazon 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.

# **Accessing Amazon Storage Gateway**

You can use the Amazon Storage Gateway console to perform various gateway configuration and management tasks. The Getting Started section and various other sections of this guide use the console to illustrate gateway functionality.

Additionally, you can use the Amazon Storage Gateway API to programmatically configure and manage your gateways. For more information about the API, see API Reference for Storage Gateway.

You can also use the Amazon SDKs to develop applications that interact with Storage Gateway. The Amazon 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 <u>Amazon</u> Developer Center.

For information about pricing, see Amazon Storage Gateway pricing.

# **Supported Amazon Regions**

- Storage Gateway For supported Amazon Regions and a list of Amazon service endpoints that you can use with Storage Gateway, see <a href="Amazon Storage Gateway endpoints">Amazon Web Services General Reference</a>.
- Storage Gateway Hardware Appliance For supported Regions that you can use with the hardware appliance, see <a href="Amazon Storage Gateway Hardware Appliance Regions">Amazon Storage Gateway Hardware Appliance Regions</a> in the Amazon Web Services General Reference.

Supported Amazon Regions API Version 2013-06-30 27

# **Using the Amazon Storage Gateway Hardware Appliance**

The Storage Gateway Hardware Appliance is a physical hardware appliance with the Storage Gateway software preinstalled on a validated server configuration. You can manage your hardware appliance from the **Hardware appliance overview** page on the Amazon 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 your hardware appliance with your Amazon account. After activation, your hardware appliance appears in the console as a gateway on the **Hardware** page. You can configure your hardware appliance as a File Gateway, Tape Gateway, or Volume Gateway type. The procedure that you use to deploy and activate these gateway types on a hardware appliance is same as on a virtual platform.

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

#### **Topics**

- Ordering information
- Supported Amazon Regions
- · Setting up your hardware appliance
- Rack-mounting your hardware appliance and connecting power
- Configuring network parameters
- Activating your hardware appliance
- Creating a gateway
- Configuring an IP address for the gateway
- Configuring your gateway
- Removing gateway software from the hardware appliance
- Deleting your hardware appliance

# **Ordering information**

The Amazon Storage Gateway hardware appliance is available exclusively through resellers. Please contact your preferred reseller for purchasing information and to request a quote.

Ordering information API Version 2013-06-30 28

# **Supported Amazon Regions**

For a list of supported Amazon Web Services Regions where the Amazon Storage Gateway Hardware Appliance is available for activation and use, see <u>Amazon Storage Gateway Hardware</u> <u>Appliance Regions</u> in the *Amazon Web Services General Reference*.

## Setting up your hardware appliance

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

#### To install and configure your hardware appliance

- Rack-mount the appliance, and plug in power and network connections. For more information, see Rack-mounting your hardware appliance and connecting power.
- 2. Set the Internet Protocol version 4 (IPv4) addresses for both the hardware appliance (the host) and Storage Gateway (the service). For more information, see Configuring network parameters.
- Activate the hardware appliance on the console Hardware appliance overview page in the Amazon Region of your choice. For more information, see Activating your hardware appliance.
- 4. Install the Storage Gateway on your hardware appliance. For more information, see Configuring 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 Amazon. 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:

Supported Amazon Regions API Version 2013-06-30 29

- 1. Reset the hardware appliance to its factory settings. Contact Amazon 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 copper network card 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

## Rack-mounting your hardware appliance and connecting power

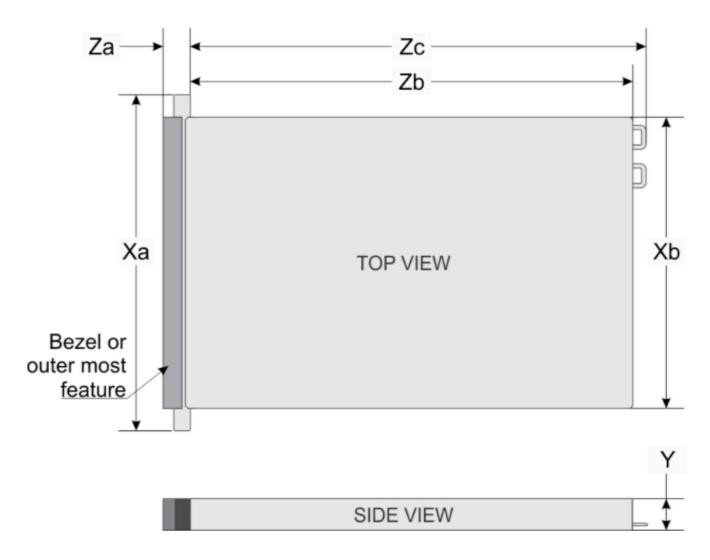
After you unbox your Storage Gateway Hardware Appliance, follow the instructions contained in the box to rack-mount the server. Your appliance has a 1U form factor and fits in a standard International Electrotechnical Commission (IEC) compliant 19-inch rack.

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.

## Hardware appliance dimensions

hardware appliance dimensions including mounting brackets and bezel.



| System              | Xa                             | Xb                             | Υ                            | Za (with<br>bezel)            | Za (without bezel)       | Zb*                             | Zc                              |
|---------------------|--------------------------------|--------------------------------|------------------------------|-------------------------------|--------------------------|---------------------------------|---------------------------------|
| 10 x 2.5-<br>inches | 482.0 mm<br>(18.97-<br>inches) | 434.0 mm<br>(17.08-<br>inches) | 42.8 mm<br>(1.68-<br>inches) | 35.84 mm<br>(1.41-<br>inches) | 22.0 mm<br>(0.87-inches) | 733.82 mm<br>(29.61-<br>inches) | 772.67 mm<br>(30.42-<br>inches) |

hardware appliance dimensions including mounting brackets and bezel.

## To connect the hardware appliance to power

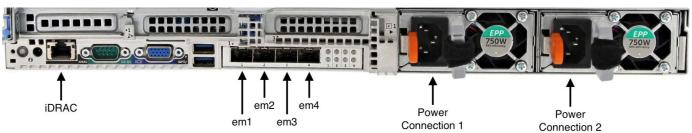


## Note

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.

Plug in a power connection to each of the two power supplies. It's possible to plug in to only 1. one power connection, but we recommend power connections to both power supplies.

In the following image, you can see the hardware appliance with the different connections. hardware appliance rear with network and power connector labels.



hardware appliance rear with network and power connector labels.

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.

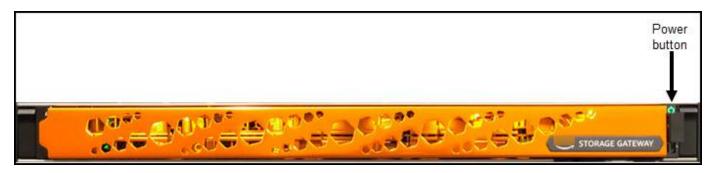


## Note

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.

- 3. Plug in the keyboard and monitor.
- Power on the server by pressing the **Power** button on the front panel, as shown in the 4. following image.

hardware appliance front with power button label.



hardware appliance front with power button label.

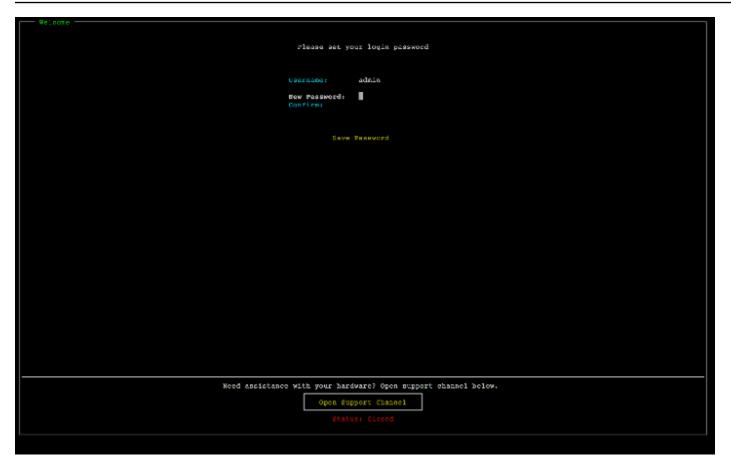
After the server boots up, the hardware console appears on the monitor. The hardware console presents a user interface specific to Amazon that you can use to configure initial network parameters. You configure these parameters to connect the appliance to Amazon and open up a support channel for troubleshooting by Amazon Support.

To work with the hardware 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.

## To set a password for the first time

- 1. For **Set Password**, enter a password, and then press Down arrow.
- 2. For **Confirm**, re-enter your password, and then choose **Save Password**.

hardware appliance console set password dialog screen.



hardware appliance console set password dialog screen.

At this point, you are in the hardware console, shown following. hardware appliance console main menu showing connections and menu options.



hardware appliance console main menu showing connections and menu options.

#### **Next step**

Configuring network parameters

# **Configuring network parameters**

After the server boots up, you can enter your first password in the hardware console as described in Rack-mounting your hardware appliance and connecting power.

Next, on the hardware console take the following steps to configure network parameters so your hardware appliance can connect to Amazon.

#### To set a network address

- Choose Configure Network and press the Enter key. The Configure Network screen shown following appears.
  - hardware appliance console configure network screen.



hardware appliance console configure network screen.

- For IP Address, enter a valid IPv4 address from one of the following sources: 2.
  - Use the IPv4 address assigned by your Dynamic Host Configuration Protocol (DHCP) server to your physical network port.

If you do so, note this IPv4 address for later use in the activation step.

• Assign a static IPv4 address. To do so, choose **Static** in the em1 section and press Enter to view the Configure Static IP screen shown following.

The em1 section is at upper left section in the group of port settings.

After you have entered a valid IPv4 address, press the Down arrow or Tab.



## Note

If you configure any other interface, it must provide the same always-on connection to the Amazon endpoints listed in the requirements.

hardware appliance console configure NIC to static IP screen.

```
Test COP. NETP

| 191 | 10.1.4.6
| Stoots 255.55.0.0
| Obtained | State |
| Stoots 255.55.0.0
| Obtained | State |
| State | State |
| Sta
```

hardware appliance console configure NIC to static IP screen.

- 3. For **Subnet**, enter a valid subnet mask, and then press Down arrow.
- 4. For **Gateway**, enter your network gateway's IPv4 address, and then press Down arrow.
- 5. For **DNS1**, enter the IPv4 address for your Domain Name Service (DNS) server, and then press Down arrow.
- (Optional) For DNS2, enter a second IPv4 address, and then press Down arrow. A second DNS server assignment would provide additional redundancy should the first DNS server become unavailable.
- 7. Choose **Save** and then press Enter to save your static IPv4 address setting for the appliance.

## To log out of the hardware console

- 1. Choose **Back** to return to the Main screen.
- 2. Choose **Logout** to return to the Login screen.

#### Next step

## Activating your hardware appliance

## **Activating your hardware appliance**

After configuring your IP address, you enter this IP address on the **Hardware** page of the Amazon Storage Gateway console to activate your hardware appliance. The activation process validates that your hardware appliance has the appropriate security credentials and registers the appliance to your Amazon account.

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

#### To activate your Amazon Storage Gateway Hardware Appliance

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



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.
- 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.

- For **Name**, enter a name for your hardware appliance. Names can be up to 255 characters long and can't include a slash character.
- 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** 

## **Creating a gateway**

You can create an S3 File Gateway, FSx File Gateway, Tape Gateway, or Volume Gateway on the hardware appliance.

## To create a gateway on your hardware appliance

- 1. Sign in to the Amazon Web Services Management Console and open the Storage Gateway console at https://console.amazonaws.cn/storagegateway/home.
- 2. Choose Hardware.
- 3. Select the activated hardware appliance on which you want to create your gateway, then choose **Create Gateway**.
- 4. Follow the procedures described in <u>Creating Your Gateway</u> to set up, connect, and configure your chosen gateway type.

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 DHCP, it can take 5–10 minutes for a gateway to display as **online** in the console. To assign a static IP address to your installed gateway, see <u>Configuring an IP address for the gateway</u>.

Creating a gateway API Version 2013-06-30 39

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 an IP address for the gateway

## Configuring an IP address for the gateway

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 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.

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

- 1. On the hardware console, choose **Open Service Console** to open a login screen for the gateway local console.
- 2. Enter the localhost **login** password, and then press Enter.
  - The default account is admin and the default password is password.
- Change the default password. Choose Actions then Set Local Password and enter your new credentials in the Set Local Password dialog box.
- 4. (Optional) Configure your proxy settings. See <u>Rack-mounting your hardware appliance and connecting power</u> for instructions.
- 5. Navigate to the Network Settings page of the gateway local console as shown following. Gateway local console configuration page showing options including network configuration.

```
Storage Gateway Configuration
## Currently connected network adapters:
##
## eth0: 10.0.0.45
1: SOCKS Proxy Configuration
2: Network Configuration
3: Test Network Connectivity
4: System Time Management
5: Gateway Console
6: View System Resource Check (0 Errors)
        Storage Gateway
0: Stop
Press "x" to exit session
Enter command: _
```

Gateway local console configuration page showing options including network configuration.

6. Type 2 to go to the **Network Configuration** page shown following.

Gateway local console network configuration page showing options including Configure Static IP.

```
Storage Gateway Network Configuration

1: Describe Adapter

2: Configure DHCP

3: Configure Static IP

4: Reset all to DHCP

5: Set Default Adapter

6: View DNS Configuration

7: View Routes

Press "x" to exit

Enter command: _____
```

Gateway local console network configuration page showing options including Configure Static IP.

7. Configure a static or DHCP IP address for the network port on your hardware appliance to present a file, volume, and Tape Gateway for applications. This IP address must be on the same subnet as the IP address used during hardware appliance activation.

## To exit the gateway local console

Press the Crt1+] (close bracket) keystroke. The hardware console appears.



## Note

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

#### **Next step**

Configuring your gateway

# **Configuring your gateway**

After your hardware appliance has been activated and configured, your appliance appears in the console. Now you can create the type of gateway that you want. Continue the installation for your gateway type. For instructions, see Configure your Amazon S3 File Gateway.

# Removing gateway software from the hardware appliance

To remove gateway software from your hardware appliance, use the following procedure.

## To remove a gateway from a hardware appliance

- On the **Hardware** page of the Storage Gateway console, choose the hardware appliance you want to delete.
- 2. For **Actions**, choose **Remove Gateway**. The confirmation dialog box appears.
- Verify that you want to remove the gateway software from specified hardware appliance, then type the word *remove* in the confirmation box and choose **Remove**.



## 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.

API Version 2013-06-30 42 Configuring your gateway

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

# **Deleting your hardware appliance**

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



#### Note

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

## To delete your hardware appliance

- If you have installed a gateway on the hardware appliance, you must first remove the gateway 1. before you can delete the appliance. For instructions on how to remove a gateway from your hardware appliance, see Removing gateway software from the hardware appliance.
- On the Hardware page of the Storage Gateway console, choose the hardware appliance you want to delete.
- For **Actions**, choose **Delete Appliance**. The confirmation dialog box appears.
- 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 topics 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:

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

## **Overview - Gateway Activation**

Gateway activation involves setting up your gateway, connecting it to Amazon, 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 onpremises 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 Amazon cloud environment. When you deploy the gateway appliance, you allocate local physical disk space on the virtualization host.

## **Connect to Amazon**

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

#### **Topics**

- Set up an Amazon S3 File Gateway
- Connect your Amazon S3 File Gateway to Amazon
- Review settings and activate your Amazon S3 File Gateway
- Configure your Amazon S3 File Gateway

Review and activate API Version 2013-06-30 45

## Set up an Amazon S3 File Gateway

#### To set up a new S3 File Gateway

- Open the Amazon Web Services Management Console at <a href="https://console.amazonaws.cn/storagegateway/home/">https://console.amazonaws.cn/storagegateway/home/</a>, and choose the Amazon Web Services Region where you want to create your gateway.
- 2. Choose **Create gateway** to open the **Set up gateway** page.
- 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 Amazon 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 Amazon 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.
- Now that your gateway is set up, you must choose how you want it to connect and communicate with Amazon. Choose **Next** to proceed.

## Connect your Amazon S3 File Gateway to Amazon

## To connect a new S3 File Gateway to Amazon

- 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 Amazon** page in the Amazon Storage Gateway console.
- 2. In the **Gateway connection options** section, for **Connection options**, choose how to identify your gateway to Amazon. You can choose from the following options:
  - IP address Provide the IP address of your gateway in the corresponding field. This IP address must be public or accessible from within your current network, and you must be able to connect to it from your web browser.
    - You can obtain the gateway IP address by logging into the gateway's local console from your hypervisor client, or by copying it from your Amazon EC2 instance details page.
  - Activation key Provide the activation key for your gateway in the corresponding field. You can generate an activation key using the gateway's local console. If your gateway's IP address is unavailable, choose this option.
- In the **Endpoint options** section, for **Service endpoint**, choose the type of endpoint that your gateway will use to communicate with Amazon. You can choose from the following options:
  - Publicly accessible Your gateway communicates with Amazon 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 Amazon 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 Amazon Regions. For more information, see Amazon Storage Gateway endpoints and quotas in the Amazon Web Services General Reference.

• VPC hosted – Your gateway communicates with Amazon 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.

Now that you have chosen how you want your gateway to connect to Amazon, 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

- 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 Amazon

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

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

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

- 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 Amazon

Review settings and activate your Amazon S3 File Gateway

When finished, choose **Next** to open the **Configure gateway** page in the Amazon 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.
- 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.
- 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.
- (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 Amazon Storage Gateway console. Repeat this step to add as many tags as you need.
- (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).

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

To check the status of your new gateway, search for it on the **Gateway overview** page of the Amazon 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 Amazon storage services without communicating over the public internet. Using the Amazon VPC service, you can launch Amazon 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 Amazon.

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.

## **Creating 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

- Sign in to the Amazon Web Services Management Console and open the Amazon VPC console 1. at https://console.amazonaws.cn/vpc/.
- In the navigation pane, choose **Endpoints**, and then choose **Create Endpoint**. 2.
- 3. On the **Create Endpoint** page, choose **Amazon Services** for **Service category**.
- 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.
- 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 look similar to this: 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.

## 

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.

## Create 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 Amazon Security Token Service (Amazon STS). If Amazon STS isn't activated in the Amazon Web Services Region where you create your File Gateway, activate it. For information about how to activate Amazon STS, see Activating and deactivating Amazon Security Token Service in an Amazon Region in the Amazon 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 with default settings
- Create an NFS file share with a custom configuration
- Create an SMB file share with default settings
- Create an SMB file share with custom configuration

# 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 Amazon 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.

Avoid unanticipated costs API Version 2013-06-30 54

# **Encrypt objects stored by File Gateway in Amazon S3**

You can use Amazon Key Management Service (Amazon KMS) to encrypt objects that your File Gateway stores in Amazon S3. To do this using the Storage Gateway console, see <u>Create an NFS file share with a custom configuration</u> or <u>Create an SMB file share with custom configuration</u>. You can also do this by using the Storage Gateway API. For instructions, see <u>CreateNFSFileShare</u> or <u>CreateSMBFileShare</u> in the *Amazon Storage Gateway API Reference*.

By default, a File Gateway uses server-side encryption managed with Amazon S3 (SSE-S3) when it writes data to an S3 bucket. If you make SSE-KMS (server-side encryption with Amazon KMS-managed keys) the default encryption for your S3 bucket, objects that a File Gateway stores there are encrypted using SSE-KMS.

To encrypt using SSE-KMS with your own Amazon KMS key, you must turn on SSE-KMS encryption. When you do so, 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 <a href="UpdateNFSFileShare"><u>UpdateNFSFileShare</u></a> or <a href="UpdateSMBFileShare"><u>UpdateSMBFileShare</u></a> API operation. This update applies to 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.

# Create an NFS file share with default settings

This section explains how to create an NFS file share using the default settings.

## **Topics**

- Create an NFS file share
- NFS file share default settings

## Create an NFS file share

Use the following procedure to quickly create a Network File System (NFS) file share with default settings.

- 1. Open the Amazon Storage Gateway console at <a href="https://console.amazonaws.cn/">https://console.amazonaws.cn/</a> storagegateway/home/ and choose **File shares** from the left navigation pane.
- 2. Choose **Create file share** to open the **File share settings** page.
- 3. For **Gateway**, choose your Amazon S3 File Gateway from the list.
- 4. For File share type, choose Network File System (NFS).
- 5. For **S3 bucket**, you can pick an existing Amazon S3 bucket in your account or create a new one using **Create new S3 bucket**. Choose the **Region** where the Amazon S3 endpoint is located and set a unique bucket name in **S3 bucket name**.

For information about creating a new bucket, see <u>How do I create an S3 bucket?</u> in the *Amazon S3 User Guide*.

If you want to connect to a bucket in another account, choose **A bucket in another account** from the drop downlist. Enter the name of the crossaccount bucket in **Cross-account bucket name**.

6. Choose **Create file share** to create an NFS file share with default settings described in the next section.

## Note

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.

## NFS file share default settings

The NFS file share that you created in the previous section uses the default settings listed in the table below.

## Note

By default, the workflow described in the previous section provides full control to the owner of the S3 bucket that maps to the NFS file share. For more information about using your file share to access objects in a bucket that's owned by another account, see <u>Using a file share for cross-account access</u>.

| Setting                       | Default value  | Notes   |
|-------------------------------|--|---|
| Amazon S3 location            | The file share is connected 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 does not include a prefix.   |
| Amazon PrivateLink for S3     | The file share is not connected to Amazon S3 through an interface endpoint in your virtual private cloud (VPC).  |   |
| File upload notification      | Off  |   |
| Storage class for new objects | 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 S3 Standard storage class, see Storage classes for frequently accessed objects in the Amazon Simple Storage Service User Guide. |
| Object metadata               | Guess MIME type  | This allows Storage Gateway to guess the MIME type for uploaded objects based on file extensions.  This option requires that Access Control Lists (ACLs) are turned on  |

| Setting                  | Default value  | Notes  |
|--------------------------|--|--|
|                          |  | for the S3 bucket associate d with your file share. If ACLs are turned off, the file share will not be able to access the Amazon S3 bucket, and will remain in the <b>Unavailable</b> state indefinitely.  |
| Enable requester pays    | Off  | For more information, see Requester Pays buckets.  |
| Audit logs               | Off  | Logging in to an Amazon<br>CloudWatch group is turned<br>off by default.   |
| 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 <a href="Supported MFS clients">Supported MFS clients</a> for File Gateway. |
| Mount options            | <ul> <li>Squash level – Root squash</li> <li>Export as – Read-write</li> </ul> | The default value of <b>Squash level</b> means that access for the remote superuser (root) is mapped to UID (65534) and GID (65534).   |

| Setting                | Default value  | Notes |
|------------------------|--|-------|
| File metadata defaults | <ul> <li>Directory permissions – 0777</li> <li>File permissions – 0666</li> <li>User ID – 65534</li> <li>Group ID – 65534</li> </ul> |       |

After your NFS file share is created, you can see your file share settings on the **Details** tab of the new file share.

Mount your NFS file share on your client

# Create an NFS file share with a custom configuration

Use the following procedure to create a Network File System (NFS) file share with custom configuration. To launch an NFS file share with default settings, see <a href="Create an NFS file share with default settings">Create an NFS file share with default settings</a>.



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

- Open the Amazon Storage Gateway console at <a href="https://console.amazonaws.cn/storagegateway/home/">https://console.amazonaws.cn/storagegateway/home/</a>.
- 2. Choose **Create file share** to open the **File share settings** page.
- 3. For **Gateway**, choose your Amazon S3 File Gateway from the list.
- 4. For **Amazon S3 location**, do one of the following:
  - To connect the file share directly to an S3 bucket, choose **S3 bucket name**, then enter the S3 bucket name and, optionally, a prefix name for objects created by the file share. Your

gateway uses this bucket to store and retrieve files. For information about creating a new bucket, see How do I create an S3 bucket? in the *Amazon S3 User Guide*.

- To connect the file share to an S3 bucket through an access point, choose S3 access point, then enter the S3 access point name and, optionally, a prefix name for objects created by the file share. Your bucket policy must be configured to delegate access control to the access point. For 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 through an access point alias, choose S3 access point alias, then enter the S3 access point alias name and, optionally, a prefix name for objects created by the file share. If you choose this option, the File Gateway cannot create a new Amazon Identity and Access Management (IAM) role and access policy on your behalf. You must select an existing IAM role and configure an access policy in the Access to your S3 bucket section that follows. 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 prefix name to prevent read/write conflicts.
- If you enter a prefix name, or choose to connect through an access point or access point alias, you must enter a file share name.
- The prefix name must end with a forward slash (/).
- After the file share is created, the prefix name can't be modified or deleted.
- For information about using prefix names, see <u>Organizing objects using prefixes</u> in the *Amazon S3 User Guide*.
- 5. For Amazon Web Services Region, choose the Amazon Web Services Region of the S3 bucket.
- 6. For **File share name**, enter a name for the file share. The default name is the S3 bucket name or access point name.

## Note

- If you entered a prefix name, or chose to connect through an access point or access point alias, you must enter a file share name.
- After the file share is created, the file share name can't be deleted.
- 7. (Optional) For **Amazon PrivateLink for S3**, do the following:
  - 1. To configure the file share to connect to S3 through an interface endpoint in your virtual private cloud (VPC) powered by Amazon PrivateLink, choose **Use VPC endpoint**.
  - 2. To identify the VPC interface endpoint that you want the file share to connect through, choose either VPC endpoint ID or VPC endpoint DNS name, and then provide the required information in the corresponding field.

## Note

- This step is required if the file share connects to S3 through a VPC access point or through an alias associated with a VPC access point.
- File share connections using Amazon PrivateLink are not supported on FIPS gateways.
- For information about Amazon PrivateLink, see <u>Amazon PrivateLink for Amazon S3</u> in the *Amazon S3 User Guide*.
- 8. For Access objects using, choose Network File System (NFS).
- 9. For **Audit logs**, choose one of the following:
  - To turn off logging, choose **Disable logging**.
  - To create a new audit log, choose **Create a new log group**.
  - To use an existing audit log, choose **Use an existing log group**, and then choose an audit log from the list.

For more information about audit logs, see <u>Understanding S3 File Gateway audit logs</u>.

10. For **Automated cache refresh from S3**, choose **Set refresh interval**, 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 the TTL interval has elapsed, accessing the directory causes the File Gateway to first refresh that directory's contents from the Amazon S3 bucket.
- 11. For File upload notification, choose Settling time (seconds) to be notified when a file has been fully uploaded to S3 by the File Gateway. Set the Settling Time in seconds to control the number of seconds to wait after the last point in time that a client wrote to a file before generating an ObjectUploaded notification. Because clients can make many small writes to files, it's best to set this parameter for as long as possible to avoid generating multiple notifications for the same file in a small time period. For more information, see Getting file upload notification.

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

- 12. (Optional) In the **Add tags** section, enter a key and value to add tags to your file share. A tag is a case-sensitive key-value pair that helps you manage, filter, and search for your file share.
- 13. Choose Next. The Configure how files are stored in Amazon S3 page appears.
- 14. For Storage class for new objects, choose a storage class to use for new objects created in your Amazon S3 bucket:
  - To store your frequently accessed object data redundantly in multiple Availability Zones that are geographically separated, choose S3 Standard. For more information about the S3 Standard storage class, see Storage classes for frequently accessed objects in the Amazon Simple Storage Service User Guide.
  - To optimize storage costs by automatically moving data to the most cost-effective storage access tier, choose S3 Intelligent-Tiering. For more information about the S3 Intelligent-Tiering storage class, see Storage class for automatically optimizing frequently and infrequently accessed objects in the Amazon Simple Storage Service User Guide.
  - To store your infrequently accessed object data redundantly in multiple Availability Zones that are geographically separated, choose S3 Standard-IA. For more information about the S3 Standard-IA storage class, see Storage classes for infrequently accessed objects in the Amazon Simple Storage Service User Guide.
  - To store your infrequently accessed object data in a single Availability Zone, choose S3 One **Zone-IA**. For more information about the S3 One Zone-IA storage class, see Storage classes for infrequently accessed objects in the Amazon Simple Storage Service User Guide.

To help monitor your S3 billing, use Amazon Trusted Advisor. For more information, see Monitoring tools in the Amazon Simple Storage Service User Guide.

- 15. For **Object metadata**, choose the metadata that you want to use:
  - To allow Storage Gateway to guess the MIME type for uploaded objects based on file extensions, choose **Guess MIME type**.
  - To give full control to the owner of the S3 bucket that maps to the NFS file share, choose **Give bucket owner full control**. For more information about using your file share to access objects in a bucket owned by another account, see Using a file share for cross-account access.

#### (i) Note

This option requires that Access Control Lists (ACLs) are turned on for the S3 bucket associated with your file share. If ACLs are turned off, the file share will not be able to access the S3 bucket, and will remain in the **Unavailable** state indefinitely.

- If you are using this file share on a bucket that requires the requester or reader instead of the bucket owner to pay for access charges, choose **Enable requester pays**. For more information, see Requester Pays buckets.
- 16. For Access to your S3 bucket, choose the Amazon Identity and Access Management (IAM) role that you want your File Gateway to use to access your Amazon S3 bucket:
  - To allow the File Gateway to create a new IAM role and access policy on your behalf, choose Create a new IAM role. This option is not available if the file share connects to Amazon S3 using an access point alias.
  - To select an existing IAM role and to set up the access policy manually, choose **Use an** existing IAM role. You must use this option if your file share connects to Amazon S3 using an access point alias. In the IAM role box, enter the Amazon Resource Name (ARN) for the role used to access your bucket. For information about IAM roles, see IAM roles in the Amazon Identity and Access Management User Guide.

For more information about access to your S3 bucket, see Granting access to an Amazon S3 bucket.

- 17. 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).
  - To use server-side encryption managed with Amazon Key Management Service (SSE-KMS), choose KMS-Managed Keys (SSE-KMS). In the Primary key box, choose an existing Amazon KMS key or choose Create a new KMS key to create a new KMS key in the Amazon Key Management Service (Amazon KMS) console. For more information about Amazon KMS, see What is Amazon Key Management Service? in the Amazon Key Management Service Developer Guide.

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

18. Choose **Next** to configure file access settings.

#### To configure file access settings

- For **Allowed clients**, specify whether to allow or restrict each client's access to your file share. Provide the IP address or CIDR notation for the clients that you want to allow. For information about supported NFS clients, see Supported NFS clients for File Gateway.
- For **Mount options**, specify the options that you want for **Squash level** and **Export as**.

For **Squash level**, choose one of the following:

- 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.
- Root squash (default): Access for the remote superuser (root) is mapped to UID (65534) and GID (65534).

For **Export as**, choose one of the following:

- Read-write
- Read-only



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

- For File metadata defaults, you can edit the Directory permissions, File permissions, User ID, and **Group ID**. For more information, see Editing the metadata defaults for your NFS file share.
- Choose Next. 4.
- Review your file share configuration settings, and then choose **Finish**.

After your NFS file share is created, you can see your file share settings in the file share's **Details** tab.

#### **Next Step**

Mount your NFS file share on your client

### Create an SMB file share with default settings

This section explains how to create an SMB file share using the default settings.



#### Note

S3 File Gateway is not designed to participate in Microsoft Distributed File System (DFS) or File Replication Service (FRS) replication, or act as a DFS namespace server. A single S3 File Gateway file share can act as a single target share within DFS, which allows it to appear in your organization's DFS listings.

#### **Topics**

- Create an SMB file share
- SMB file share default settings

## Create an SMB file share

### **SMB File share prerequisite**

Before you create a Server Message Block (SMB) file share, make sure that you configure SMB security settings for your File Gateway. You must also configure either Microsoft Active Directory (AD) or guest access for authentication. A file share provides one type of SMB access only. For instructions, see Editing SMB settings for a gateway.

### SMB file share creation with default settings

Use the following procedure to create an SMB file share using the default settings.



#### Note

An SMB file share doesn't work properly unless the required ports are open in your security group. For more information, see Port Requirements.

- Open the Amazon Storage Gateway console at https://console.amazonaws.cn/ 1. storagegateway/home/ and choose **File shares** from the left navigation pane.
- 2. Choose **Create file share** to open the **File share settings** page.
- 3. For **Gateway**, select your Amazon S3 File Gateway from the list.
- 4. For **File share type**, choose **Server Message Block (SMB)**.
- For S3 bucket, you can select an existing Amazon S3 bucket in your account or create a new 5. one using Create new S3 bucket. Choose the Region where the Amazon S3 endpoint is located and set a unique bucket name in S3 bucket name.

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

If you want to connect to a bucket in another account, choose A bucket in another account from the dropdown list. Enter the name of the account bucket in **Cross-account bucket name**.

- Choose **Configure** to authenticate the configuration.
  - For **Authentication method**, choose the authentication method that you want to use.
    - To use your corporate Microsoft AD for user authenticated access to your SMB file share, choose Active Directory.

Create an SMB file share API Version 2013-06-30 66

Your File Gateway must be joined to a domain. For more information, see Using Active Directory to authenticate users.

Joining a domain creates an AD 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).

If your AD 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 AD environment has a designated organizational unit (OU) for new computer objects, you must specify that OU when joining the domain.

- To provide only guest access, choose **Guest access**. If you choose this authentication method, your File Gateway doesn't have to be part of a Microsoft AD domain. You can also use a File Gateway that's a member of an AD domain to create file shares with guest access. You must set a guest password for your SMB server in the corresponding field.
- 7. Choose Create file share to create an SMB file share using the default settings described in the next section.

#### Note

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.

### SMB file share default settings

The SMB file share that you created in the previous section uses the default settings listed in the table below.



#### Note

By default, the work flow described in the previous section provides full control to the owner of the Amazon S3 bucket that maps to the SMB file share. For more information about using your file share to access objects in a bucket owned by another account, see <u>Using a file share for cross-account access</u>.

| Setting                       | Default value  | Notes   |
|-------------------------------|--|---|
| Amazon S3 location            | The file share is connected 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.  |
| Amazon PrivateLink for S3     | The file share isn't connected to Amazon S3 through an interface endpoint in your virtual private cloud (VPC).   |   |
| File upload notification      | Off  |   |
| Storage class for new objects | 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 S3 Standard storage class, see Storage classes for frequently accessed objects in the Amazon Simple Storage Service User Guide. |
| Object metadata               | Guess MIME type  | This allows Storage Gateway to guess the MIME type for uploaded objects based on file extensions.   |

| Setting                  | Default value | Notes   |
|--------------------------|---------------|---|
|                          |               | This option requires that Access Control Lists (ACLs) are turned on for the S3 bucket that's associated with your file share. If ACLs are turned off, the file share won't be able to access the Amazon S3 bucket, and will remain in the Unavailable state indefinitely. |
| Access based enumeration | Not activated | The files and folders on the file share are visible to all users during directory enumerati on. Access-based enumerati on 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).            |
| Enable requester pays    | Off           | For more information, see Requester Pays buckets.   |

| Setting                     | Default value         | Notes  |
|-----------------------------|-----------------------|--|
| Opportunistic lock (oplock) | On                    | This allows the file share to use opportunistic locking to optimize the file buffering strategy. In most cases, activatin g opportunistic locking improves performance, particularly with regard to Windows context menus. |
| Audit logs                  | Off                   | Logging in to an Amazon<br>CloudWatch group is turned<br>off by default.   |
| Force case sensitivity      | On                    | 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.  |

After your SMB file share is created, you can see your file share settings on the **Details** tab of the new file share.

Mount your SMB file share on your client

## Create an SMB file share with custom configuration

Before you create a Server Message Block (SMB) file share, make sure that you configure SMB security settings for your File Gateway. You must also configure either Microsoft Active Directory

(AD) or guest access for authentication. A file share provides one type of SMB access only. For instructions, see Editing SMB settings for a gateway.

#### Note

- An SMB file share does not work properly unless the required ports are open in your security group. For more information, see <u>Port Requirements</u>.
- S3 File Gateway is not designed to participate in Microsoft Distributed File System (DFS) or File Replication Service (FRS) replication, or act as a DFS namespace server. A single S3 File Gateway file share can act as a single target share within DFS, which allows it to appear in your organization's DFS listings.
- 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.

### Creating an SMB file share

#### To create an SMB file share

- Open the Amazon Storage Gateway console at <a href="https://console.amazonaws.cn/storagegateway/home/">https://console.amazonaws.cn/storagegateway/home/</a>.
- 2. Choose **Create file share** to open the **File share settings** page.
- 3. For **Gateway**, choose your Amazon S3 File Gateway from the list.
- 4. For **Amazon S3 location**, do one of the following:
  - To connect the file share directly to an S3 bucket, choose **S3 bucket name**, then enter the bucket name and, optionally, a prefix name for objects created by the file share. Your gateway uses this bucket to store and retrieve files. For information about creating a new bucket, see How do I create an S3 bucket? in the *Amazon S3 User Guide*.
  - To connect the file share to an S3 bucket through an access point, choose **S3 access point**, then enter the S3 access point name and, optionally, a prefix name for objects created by the file share. Your bucket policy must be configured to delegate access control to the access point. For information about access points, see <a href="Managing data access with Amazon S3 access">Managing data access with Amazon S3 access</a> points and Delegating access control to access points in the Amazon S3 User Guide.

To connect the file share to an S3 bucket through an access point alias, choose S3 access point alias, then enter the S3 access point alias name and, optionally, a prefix name for objects created by the file share. If you choose this option, the File Gateway cannot create a new Amazon Identity and Access Management (IAM) role and access policy on your behalf. You must select an existing IAM role and configure an access policy in the Access to your S3 bucket section that follows. 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 prefix
  name to prevent read/write conflicts.
- If you enter a prefix name, or choose to connect through an access point or access point alias, you must enter a file share name.
- The prefix name must end with a forward slash (/).
- After the file share is created, the prefix name can't be modified or deleted.
- For information about using prefix names, see <u>Organizing objects using prefixes</u> in the *Amazon S3 User Guide*.
- 5. For **Amazon Web Services Region**, choose the Amazon Web Services Region of the S3 bucket.
- 6. For **File share name**, enter a name for the file share. The default name is the S3 bucket name or access point name.

### Note

- If you entered a prefix name, or chose to connect through an access point or access point alias, you must enter a file share name.
- After the file share is created, the file share name can't be deleted.
- 7. (Optional) For Amazon PrivateLink for S3, do the following:
  - 1. To configure the file share to connect to S3 through an interface endpoint in your virtual private cloud (VPC) powered by Amazon PrivateLink, choose **Use VPC endpoint**.

Creating an SMB file share API Version 2013-06-30 72

2. To identify the VPC interface endpoint that you want the file share to connect through, choose either VPC endpoint ID or VPC endpoint DNS name, and then provide the required information in the corresponding field.

### Note

- This step is required if the file share connects to S3 through a VPC access point or through an alias associated with a VPC access point.
- File share connections using Amazon PrivateLink are not supported on FIPS gateways.
- For information about Amazon PrivateLink, see <u>Amazon PrivateLink for Amazon S3</u> in the *Amazon Simple Storage Service User Guide*.
- 8. For Access objects using, choose Server Message Block (SMB).
- 9. For **Audit logs**, choose one of the following:
  - To turn off logging, choose **Disable logging**.
  - To create a new audit log, choose **Create a new log group**.
  - To use an existing log group, choose **Use an existing log group**, and then choose an audit log from the list.

For more information about audit logs, see Understanding S3 File Gateway audit logs.

- 10. For **Automated cache refresh from S3**, choose **Set refresh interval**, and then 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 the TTL interval has elapsed, accessing the directory causes the File Gateway to first refresh that directory's contents from the Amazon S3 bucket.
- 11. For **File upload notification**, choose **Settling time** (seconds) to be notified when a file has been fully uploaded to S3 by the File Gateway. Set the **Settling Time** in seconds to control the number of seconds to wait after the last point in time that a client wrote to a file before generating an ObjectUploaded notification. Because clients can make many small writes to files, it's best to set this parameter for as long as possible to avoid generating multiple notifications for the same file in a small time period. For more information, see <u>Getting file</u> upload notification.



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

- 12. (Optional) In the Tags section, choose Add new tag, and then enter a key and value to add tags to your file share. A tag is a case-sensitive key-value pair that helps you manage, filter, and search for your file share.
- 13. Choose **Next**. The **Amazon S3 storage settings** page appears.
- 14. For **Storage class for new objects**, choose a storage class to use for new objects created in your Amazon S3 bucket:
  - To store your frequently accessed object data redundantly in multiple Availability Zones that are geographically separated, choose S3 Standard. For more information about the S3 Standard storage class, see Storage classes for frequently accessed objects in the Amazon Simple Storage Service User Guide.
  - To optimize storage costs by automatically moving data to the most cost-effective storage access tier, choose S3 Intelligent-Tiering. For more information about the S3 Intelligent-Tiering storage class, see Storage class for automatically optimizing frequently and infrequently accessed objects in the Amazon Simple Storage Service User Guide.
  - To store your infrequently accessed object data redundantly in multiple Availability Zones that are geographically separated, choose S3 Standard-IA. For more information about the S3 Standard-IA storage class, see Storage classes for infrequently accessed objects in the Amazon Simple Storage Service User Guide.
  - To store your infrequently accessed object data in a single Availability Zone, choose S3 One **Zone-IA**. For more information about the S3 One Zone-IA storage class, see Storage classes for infrequently accessed objects in the Amazon Simple Storage Service User Guide.

To help monitor your S3 billing, use Amazon Trusted Advisor. For more information, see Monitoring tools in the Amazon Simple Storage Service User Guide.

- 15. For **Object metadata**, choose the metadata that you want to use:
  - To allow Storage Gateway to guess the MIME type for uploaded objects based on file extensions, choose Guess MIME type.

• To give full control to the owner of the S3 bucket that maps to the SMB file share, choose **Give bucket owner full control**. For more information about using your file share to access objects in a bucket owned by another account, see Using a file share for cross-account access.

#### Note

This option requires that Access Control Lists (ACLs) are turned on for the S3 bucket associated with your file share. If ACLs are turned off, the file share will not be able to access the S3 bucket, and will remain in the **Unavailable** state indefinitely.

- To give full control to the owner of the S3 bucket that maps to the SMB file share, choose **Enable requester pays.** For more information, see Requester Pays buckets.
- 16. For Access to your S3 bucket, choose the Amazon Identity and Access Management (IAM) role that you want your File Gateway to use to access your Amazon S3 bucket:
  - To allow the File Gateway to create a new IAM role and access policy on your behalf, choose Create a new IAM role. This option is not available if the file share connects to Amazon S3 using an access point alias.
  - To select an existing IAM role and to set up the access policy manually, choose **Use an** existing IAM role. You must use this option if your file share connects to Amazon S3 using an access point alias. In the IAM role box, enter the Amazon Resource Name (ARN) for the role used to access your bucket. For information about IAM roles, see IAM roles in the Amazon Identity and Access Management User Guide.

For more information about access to your S3 bucket, see Granting access to an Amazon S3 bucket.

- 17. 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).
  - To use server-side encryption managed with Amazon Key Management Service (SSE-KMS), choose **KMS-Managed Keys (SSE-KMS)**. In the **Primary key** box, choose an existing Amazon KMS key or choose **Create a new KMS key** to create a new KMS key in the Amazon Key Management Service (Amazon KMS) console. For more information about Amazon KMS,

see What is Amazon Key Management Service? in the Amazon Key Management Service Developer Guide.



#### Note

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

- 18. Choose **Next**. The **File access settings** page appears.
- 19. For **Authentication method**, choose the authentication method that you want to use.
  - To use your corporate Microsoft AD for user authenticated access to your SMB file share, choose **Active Directory**.



#### Note

Your File Gateway must be joined to a domain. For more information, see Using Active Directory to authenticate users.

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).

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.

• To provide only guest access, choose **Guest access**. If you choose this authentication method, your File Gateway doesn't have to be part of a Microsoft AD domain. You can also use a File Gateway that is a member of an AD domain to create file shares with guest access. You must set a guest password for your SMB server in the corresponding field.

Creating an SMB file share API Version 2013-06-30 76



Both access types are available at the same time.

20. In the **SMB share settings** section, choose your settings.

For **Export as**, choose one of the following:

- **Read-write** (the default value)
- Read-only



#### (i) Note

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

For **File/directory access controlled by**, choose one of the following:

- To set fine-grained permissions on files and folders in your SMB file share, choose **Windows** Access Control List. For more information, see Using Windows ACLs to limit SMB file share access.
- To use POSIX permissions to control access to files and directories that are stored through an NFS or SMB file share, choose **POSIX permissions**.

If your authentication method is **Active Directory**, for **Admin users/groups**, enter a commaseparated list of AD users and groups. Do this if you want the admin user to have privileges to update access control lists (ACLs) on all files and folders in the file share. These users and groups then have administrator rights to the file share. A group must be prefixed with the @ character, for example, @group1.

For **Opportunistic lock (oplock)**, choose one of the following:

- To allow the file share to use opportunistic locking to optimize the file buffering strategy, select **Activate**. In most cases, activating opportunistic locking improves performance, particularly with regard to Windows context menus.
- To prevent the use of opportunistic locking, deselect Activate. If multiple Windows clients in your environment frequently edit the same files simultaneously, deactivating opportunistic locking can sometimes improve performance.



To avoid conflict cases during file access, this setting cannot be configured if **Force** case sensitivity is selected.

#### For **Force case sensitivity**, choose one of the following:

- To allow the client to control the case sensitivity, select Activate.
- To allow the gateway to control the case sensitivity, deselect **Activate**.



#### Note

When activated, this setting applies immediately to new SMB client connections. Existing SMB client connections must disconnect from the file share and reconnect for the setting to take effect.

To avoid conflict cases during file access, this setting cannot be configured if **Opportunistic lock** is selected.

#### For **Access based enumeration**, choose one of the following:

- To make the files and folders on the share visible only to users who have read access, select Activate.
- To make the files and folders on the share visible to all users during directory enumeration, deselect Activate.



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).

21. (Optional) 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.

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.

- 22. Choose Next.
- 23. Review your file share configuration settings, and then choose **Finish**.

After your SMB file share is created, you can see your file share settings in the file share's Details tab.

#### **Next Step**

Mount your SMB file share on your client

## Mount and use your file share

Following, you can find instructions about how to mount your file share on your client, use your share, test your File Gateway, and clean up resources as needed. For more information about supported Network File System (NFS) clients, see <a href="Supported NFS clients for File Gateway">Supported NFS clients for File Gateway</a>. For more information about supported Service Message Block (SMB) clients, see <a href="Supported SMB clients for File Gateway">Supported SMB clients for File Gateway</a>.

You can find example commands to mount your file share on the Amazon Web Services Management Console. In following sections, you can find details on how to mount your file share on your client, use your share, test your File Gateway, and clean up resources as needed.

#### **Topics**

- Mount your NFS file share on your client
- Mount your SMB file share on your client
- Using file shares on buckets with pre-existing objects
- Test your S3 File Gateway
- Cleaning up resources you don't need

### Mount your NFS file share on your client

Now you 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 <u>create an SMB file share</u> 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.
- 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 MacOS clients, type the following command at the command prompt.

sudo mount -t nfs -o vers=3.0,rsize=1048576,wsize=1048576,hard,nolock -v [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 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:

- 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.
- By default Windows uses a soft mount for mounting your NFS share. Soft mounts time out more easily when there are connection issues. We recommend using a hard mount because a hard mount is safer and better preserves your data. The soft mount command omits the **-o mtype=hard** switch. The Windows hard mount command uses the -o mtype=hard switch.
- If you are using Windows clients, check your mount options after mounting by running the mount command with no options. The response should that 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

Now you 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:

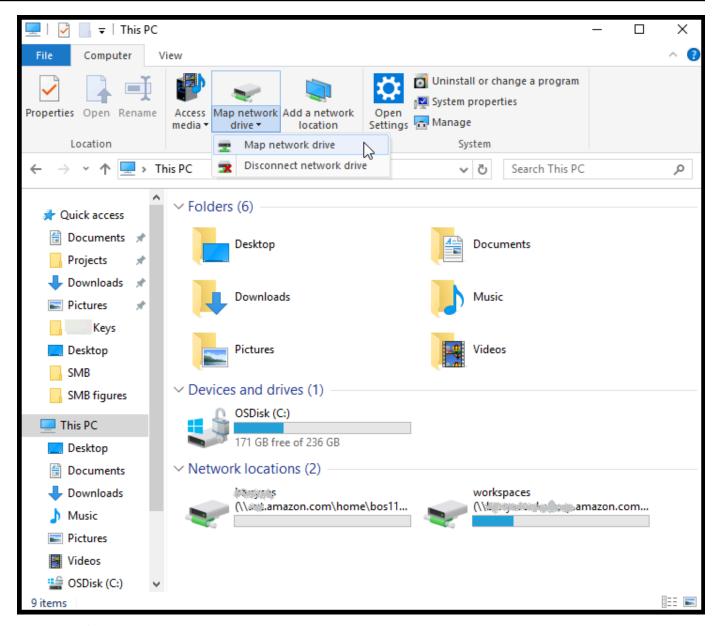
 You might have a case where a folder and an object exist in an Amazon S3 bucket and have the same name. In this case, if the object name doesn't contain a trailing slash, only

- 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 choose **Map Network Drive** for **Map Network Drive** in the **Computer** tab, as shown in the following screen shot.

  windows file explorer with this pc, the computer tab, and map network drive selected.



windows file explorer with this pc, the computer tab, and map network drive selected.

- 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.

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.amazonaws.cn/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.
- For Actions, choose the action that you want: 4.
  - 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.

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 "Refresh Amazon S3 bucket objects".



#### (i) 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**

You can copy files and folders to your mapped drive. The files automatically upload to your Amazon S3 bucket.

#### To upload files from your Windows client to Amazon S3

On your Windows client, navigate to the drive that you mounted your file share on. The name 1. of your drive is preceded by the name of your S3 bucket.

- Copy files or a folder to the drive. 2.
- 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 Amazon 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 1s -1) depends on the number of objects in your bucket.

#### **Next Step**

Cleaning up resources you don't need

### Cleaning up resources you don't need

If you created your gateway as an example exercise or a test, consider cleaning up to avoid incurring unexpected or unnecessary charges.

#### To clean up resources you don't need

- 1. Unless you plan to continue using the gateway, delete it. For more information, see <u>Deleting</u> 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.

## Managing your Amazon S3 File Gateway

Following, you can find information about how to manage your Amazon S3 File Gateway resources.

#### **Topics**

- Editing Basic Gateway Information
- Granting access and permissions for file shares and buckets
- Deleting a file share
- Editing SMB settings for a gateway
- · Editing settings for your NFS file share
- Editing the metadata defaults for your NFS file share
- Editing access settings for your NFS file share
- Editing settings for your SMB file share
- Editing access settings for your SMB file share
- Refreshing objects in your Amazon S3 bucket
- Using S3 Object Lock with an Amazon S3 File Gateway
- Understanding file share status
- Understanding gateway status

## **Editing Basic Gateway Information**

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 <a href="https://console.amazonaws.cn/storagegateway/home">https://console.amazonaws.cn/storagegateway/home</a>.
- 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. Modify the settings you want to change, then choose **Save changes**.



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.

### 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. Buckets that you can grant access to include buckets in a different Amazon Web Services account than your file share. For information about how to add a file share, see Create a file share.

#### **Topics**

- Granting access to an Amazon S3 bucket
- Cross-service confused deputy prevention
- Using a file share for cross-account access

### 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 Amazon 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 Amazon Web Service in the IAM User Guide.

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

#### ▲ 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 Amazon Identity and Access Management User Guide:

- IAM: Pass an IAM role to a specific Amazon service
- Granting a user permissions to pass a role to an Amazon 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 <a href="Create a file share">Create a file share</a>.

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 TestBucket. The second part allows the listed actions on all objects in TestBucket.

```
{
    "Version": "2012-10-17",
    "Statement": [
```

```
{
            "Action": [
                "s3:GetAccelerateConfiguration",
                "s3:GetBucketLocation",
                "s3:GetBucketVersioning",
                "s3:ListBucket",
                "s3:ListBucketVersions",
                "s3:ListBucketMultipartUploads"
            ],
            "Resource": "arn:aws:s3:::TestBucket",
            "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:::TestBucket/*",
            "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.

If you need to connect your file share to an S3 bucket through a VPC endpoint, see Endpoint policies for Amazon S3 in the Amazon 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 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 <u>Using Identity-Based Policies</u> (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 Amazon, 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, Amazon 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 <a href="mailto:aws:SourceArn">aws:SourceAccount</a> global condition context keys in resource policies to limit the permissions that Amazon 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.

```
{
  "Version": "2012-10-17",
  "Statement": {
    "Sid": "ConfusedDeputyPreventionExamplePolicy",
      "Effect": "Allow",
      "Principal": {
        "Service": "storagegateway.amazonaws.com"
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "123456789012"
        },
        "ArnLike": {
          "aws:SourceArn": "arn:aws:storagegateway:us-east-1:123456789012:gateway/
sgw-712345DA"
        }
      }
    }
  ]
```

}

## 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 <a href="Example 2: Bucket owner granting cross-account bucket permissions">Example 2: Bucket owner granting cross-account bucket permissions</a> 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 acccessible to S3 bucket owner** when creating your file share or editing file share settings in the https://console.amazonaws.cn/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

## Deleting 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.amazonaws.cn/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.
- Verify that you want to delete the specified file shares, then type the word delete in the 4. 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

Deleting a file share API Version 2013-06-30 97 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 Amazon Web Services 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

- Open the Storage Gateway console at https://console.amazonaws.cn/storagegateway/home. 1.
- From the **File shares** list page, choose file share that you flagged for deletion in the procedure 2. above to view its details. After a few seconds, a deletion notification message appears on the **Details** tab.
- 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 multipart 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

Open the Storage Gateway console at https://console.amazonaws.cn/storagegateway/home. 1.

- Choose **Gateways**, then choose the gateway for which you want to edit SMB settings. 2.
- From the Actions dropdown menu, choose Edit SMB settings, then choose the settings you 3. want to edit.

See the following topics for more information.

#### **Topics**

- Setting a security level for your gateway
- Using Active Directory to authenticate users
- Providing guest access to your file share
- Configure Local Groups for your gateway
- Setting file share visibility

# Setting 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

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



#### Note

This setting is called SMBSecurityStrategy in the API Reference. A higher security level can affect performance.

• Enforce encryption – If you choose this option, S3 File Gateway only allows connections from SMBv3 clients that have encryption turned on. This option is highly recommended for environments that handle sensitive data. This option works with SMB clients on Microsoft Windows 8, Windows Server 2012, or later.

- 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 SMB clients on Microsoft Windows Vista, Windows Server 2008, or later.
- 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.

Choose Save. 5.

# Using Active Directory to authenticate users

To use your corporate Active Directory 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 Amazon Directory Service, you can create a hosted Active Directory domain service in the Amazon Web Services Cloud.

Anyone who can provide the correct password gets guest access to the SMB file share.

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

- Open the Storage Gateway console at https://console.amazonaws.cn/storagegateway/home. 1.
- 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.
- 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).

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.

- For **Domain user** and **Domain password**, enter the credentials for the Active Directory service account that the gateway will use to join the domain.
- (Optional) For Organization unit (OU), enter the designated OU that your Active Directory uses for new computer objects.
- (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.
- 8. Choose Save changes.

#### To limit file share access to specific AD users and groups

In the Storage Gateway console, choose the file share that you want to limit access to. 1.

- From the Actions drop-down menu, choose Edit file share access settings. 2.
- 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.

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

# Providing guest access to your file share

If you want to provide only guest access, your S3 File Gateway doesn't have to be part of a Microsoft AD domain. You can also use a S3 File Gateway that is a member of an AD domain to create file shares with guest access. Before you create a file share using guest access, you need to change the default password.

These steps set the password for the quest user smbquest. The smbquest user is the user when the authentication method for the file share is set to GuestAccess.

#### To change the guest access password

- Open the Storage Gateway console at https://console.amazonaws.cn/storagegateway/home. 1.
- 2. Choose **Gateways**, then choose the gateway for which you want to edit SMB settings.
- From the Actions drop-down menu, choose Edit SMB settings, then choose Guest access 3. settings.
- For **Guest password**, provide a password, and then choose **Save**.

# **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

- Open the Storage Gateway console at https://console.amazonaws.cn/storagegateway/home. 1.
- Choose **Gateways**, then choose the gateway for which you want to edit SMB settings. 2.
- 3. From the Actions dropdown menu, choose Edit SMB settings, then choose Local Group settings.
- 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.

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

# Setting file share visibility

File share visibility controls whether the shares on a gateway are visible when listing shares to users.

#### To set file share visibility

- Open the Storage Gateway console at https://console.amazonaws.cn/storagegateway/home. 1.
- 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.
- For **Visibility status**, select the check box to have the shares on this gateway appear when listing shares to users. Keep the check box cleared to have the shares on this gateway not appear when listing shares to users.

# **Editing settings for your NFS file share**

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

- File share name
- Audit logs
- Existing log group list
- Non-gateway file cache refresh time
- Upload events settling time

Setting file share visibility API Version 2013-06-30 104

- Storage class for new objects
- Guess MIME type
- Gateway files acccessible to S3 bucket owner
- Enable requester pays
- Access level
- Export as

#### 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

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

# Editing the metadata defaults for your NFS file share

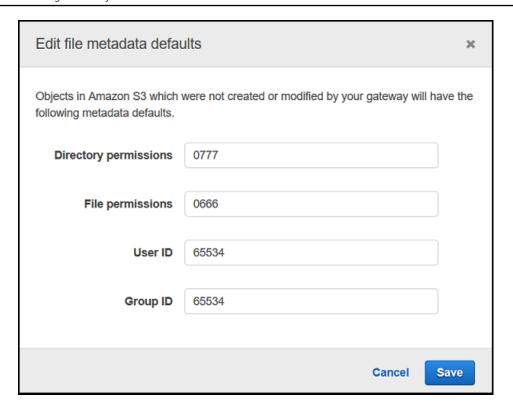
If you don't set metadata values for your files or directories in your bucket, your S3 File Gateway sets default metadata values. These values include Unix permissions for files and folders. You can edit the metadata defaults on the Storage Gateway console.

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.                            |
| User ID               | The default owner ID for files in the file share.<br>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.amazonaws.cn/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**.



# **Editing access settings for your NFS file share**

We recommend changing the allowed NFS client settings for your NFS file share. If you don't, any client on your network can mount to your file share.

#### To edit NFS access settings

- 1. Open the Storage Gateway console at https://console.amazonaws.cn/storagegateway/home.
- 2. Choose File shares, and then choose the NFS file share that you want to edit.
- 3. For Actions, choose Edit share access settings.
- 4. In the **Edit allowed clients** dialog box, choose **Add entry**, provide the IP address or CIDR notation for the clients that you want to allow, and then choose **Save**.

# **Editing 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
- **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 acccessible to S3 bucket owner** choose whether to make files on the gateway accessible to the Amazon 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 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.amazonaws.cn/storagegateway/home.

- Choose File shares, and then choose the file share that you want to update. 2.
- 3. For Actions, choose Edit file share settings.
- Edit any settings that you want to change. 4.
- Choose Save. 5.

# Editing access settings 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.amazonaws.cn/storagegateway/home.
- 2. Choose **File shares**, and then choose the SMB file share that you want to edit.
- For **Actions**, choose **Edit file share access settings**. 3.
- In the **User and group file share access** section, choose your settings. 4.

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.

# Refreshing objects in your Amazon S3 bucket

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.

#### **Topics**

- Configure an automated cache refresh schedule using the Storage Gateway console
- Configure an automated cache refresh schedule using Amazon 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

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

- 1. Open the Storage Gateway console at <a href="https://console.amazonaws.cn/storagegateway/home">https://console.amazonaws.cn/storagegateway/home</a>.
- 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 Amazon Lambda with an Amazon CloudWatch rule

# To configure an automated cache refresh schedule using Amazon 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.

```
{
    "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):
```

- 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 Amazon CLI.

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

The Lambda invocation gives you output similar to the following.

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 Amazon CLI.
- 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

- Open the Storage Gateway console at https://console.amazonaws.cn/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

#### 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 Amazon 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 an 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 <u>Locking objects using S3 Object Lock</u> 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 Amazon.                                      |
| 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 Amazon.  |
| DELETING       | The file share is being deleted. The file share is not deleted until all data is uploaded to Amazon. 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 Amazon. The FORCE_DEL ETING 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 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 gateway status on the Storage Gateway console. Gateway status appears in the **Status** column for each of your gateways. 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 | <ul> <li>Your gateway might be in an OFFLINE status for one or more of the following reasons:</li> <li>The gateway can't reach the Storage Gateway service endpoints.</li> <li>The gateway had an unexpected shutdown.</li> <li>The gateway has an associated cache disk that is disconnected, has been modified, or has failed.</li> </ul> |

# **Monitoring Storage Gateway**

This section describes 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 Amazon 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 <a href="mailto:Amazon CloudWatch User Guide">Amazon CloudWatch User Guide</a>. For more information about CloudWatch, see the Amazon CloudWatch User Guide.

#### **Topics**

- Understanding CloudWatch alarms
- Creating recommended CloudWatch alarms for your gateway
- Creating a custom CloudWatch alarm for your gateway
- Monitoring your S3 File 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 <a href="Understanding gateway metrics">Understanding gateway metrics</a> and <a href="Understanding file share metrics">Understanding file share metrics</a>. 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 <a href="Using Amazon CloudWatch alarms">Using Amazon CloudWatch alarms</a> 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   |
|-------|---|
| ОК    | The metric or expression is within the defined threshold.     |
| Alarm | The metric or expression is outside of the defined threshold. |

| State             | Description  |
|-------------------|--|
| 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 <a href="Creating a custom">Creating a custom</a> <a href="CloudWatch alarm">CloudWatch alarm for your gateway</a> . |
| Unavailable       | The state of the alarm is unknown. Choose <b>Unavailable</b> to view error information in the <b>Monitoring</b> 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 <a href="Configure your Amazon S3 File Gateway">Configure your Amazon S3 File Gateway</a>. If you want to add or update recommended CloudWatch alarms for an existing gateway, use the following procedure.

#### To add or update recommended CloudWatch alarms for an existing gateway



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.amazonaws.cn/storagegateway/home/.

- 2. In the navigation pane, choose **Gateways**, then choose the gateway for which you want to create recommended CloudWatch alarms.
- 3. On the gateway details page, 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.

# Creating 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.amazonaws.cn/storagegateway/home/.
- 2. In the navigation pane, choose **Gateways**, then choose the gateway for which you want to create an alarm.
- 3. On the gateway details page, choose the **Monitoring** tab.
- 4. Under **Alarms**, choose **Create alarm** to open the CloudWatch console.
- 5. Use the CloudWatch console to create the type of alarm that you want. You can create the following types of alarms:
  - Static threshold alarm: An alarm based on a set threshold for a chosen metric. The alarm enter the ALARM state when the metric breaches the threshold for a specified number of evaluation periods.

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

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

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

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

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

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

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

- 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 Amazon 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 <a href="Real-time Processing of Log Data with">Real-time Processing of Log Data with</a> 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 <a href="Configure your Amazon S3 File Gateway">Configure your Amazon S3 File Gateway</a>. For general information about CloudWatch log groups, see <a href="Working with Log Groups and Log Streams">Working with Log Groups and Log Streams</a> 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/my-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 Amazon Web Services Management Console and open the Storage Gateway console at https://console.amazonaws.cn/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**. Or, on the **Details** tab, under **Health logs** and **Not Enabled**, choose **Configure log group** to open the **Edit** *CustomerGatewayName* dialog box.
- 4. For **Gateway health log group**, choose one of the following:
  - Disable logging if you don't want to monitor your gateway using CloudWatch log groups.
  - 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**.

- 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.

#### To configure a CloudWatch Log Group to work with your S3 File Gateway

- 1. Sign in to the Amazon Web Services Management Console and open the Storage Gateway console at <a href="https://console.amazonaws.cn/storagegateway/home">https://console.amazonaws.cn/storagegateway/home</a>.
- 2. Choose **Gateways**, and then choose the gateway that you want to configure the CloudWatch log group for.
- 3. For **Actions**, choose **Edit gateway information**. Or, in the **Details** tab, next to **Logging**, under **Not Enabled**, choose **Configure log group** to open the **Edit gateway information** dialog box.
- 4. For **Gateway log group**, choose **Use an existing log group**, and then choose the log group that you want to use.
  - If you don't have a log group, choose **Create a new log group** to create one. You are directed to the CloudWatch Logs console where you can create the log group. If you create a new log group, choose the refresh button to view the new log group in the drop-down list.
- 5. When you are done, choose **Save**.
- 6. To see the logs for your gateway, choose the gateway, and then choose the **Details** tab.

For information about how to troubleshoot errors, see Troubleshooting: File Gateway issues.

# **Using Amazon CloudWatch metrics**

You can get monitoring data for your S3 File Gateway by using either the Amazon Web Services 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 <a href="Management Console">Amazon CloudWatch API</a> 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 <a href="Dimensions">Dimensions</a> 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<br>CloudWatch<br>namespace | Dimension                  | Description   |
|-----------------------------------|----------------------------|---|
| AWS/Stora<br>geGateway            | GatewayId ,<br>GatewayName | 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 GatewayId and the GatewayName dimensions.  Throughput and latency data of a gateway are based on all the file shares in the gateway.  Data is available automatically in 5-minute periods at no charge. |

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

- Viewing available metrics
- Getting statistics for a metric
- Creating CloudWatch alarms

# **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 <u>NotifyWhenUploaded</u> 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 <u>Getting working file set upload</u> notification.
- You can get notified when the gateway finishes refreshing the cache for your S3 bucket. When you invoke the <u>RefreshCache</u> operation through the Storage Gateway console or API, subscribe to the notification when the operation is complete. For more information, see <u>Getting refresh</u> 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 Amazon 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 Amazon Region. For example, you could specify the Storage Gateway File Upload Event as the event type.

For information about how to use CloudWatch Events rules, see <u>Creating a CloudWatch Events rule</u> 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 Amazon 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 <u>RefreshCache</u> request for the file share on Gateway2. When the <u>RefreshCache</u> 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": "my-bucket",
    }
}
```

| 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 Amazon service that is the source of the request and notification.               |
| account           | The ID of the Amazon account where the request and notification were generated from. |
| time              | When the request to upload files to Amazon S3 was made.                              |
| region            | The Amazon 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 Amazon 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 <a href="NotifyWhenUploaded">NotifyWhenUploaded</a> 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 <a href="RefreshCache">RefreshCache</a> request for the file share on Gateway2. When the <a href="RefreshCache">RefreshCache</a> 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 Upload Notification Event.

```
"version": "2012-10-17",
    "id": "2649b160-d59d-c97f-3f64-8aaa9ea6aed3",
    "detail-type": "Storage Gateway Upload Notification 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",
```

| 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 Amazon service that is the source of the request and notification.  |
| account          | The ID of the Amazon account where the request and notification were generated from.  |
| time             | When the request to upload files to Amazon S3 was made.   |
| region           | The Amazon 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 NotifyWhenUploaded is called. |
| request-received | When the gateway received the NotifyWhe nUploaded request.  |

| Field names | Description   |
|-------------|---|
| 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",
```

| 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 Amazon service that is the source of the request and notification.  |
| account         | The ID of the Amazon account where the request and notification were generated from.  |
| time            | When the request to refresh the files in working-set was made.  |
| region          | The Amazon 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 RefreshCache . |
| started         | when the gateway received the RefreshCa che request and the refresh was started.  |

| Field names | Description  |
|-------------|--|
| completed   | When the refresh of the working-set was completed.   |
| 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 fileshare-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  |
|---------------------------|--|
| AvailabilityNotifications | This metric reports the number of availabil ity-related health notifications that were generated by the gateway in the reporting period. |

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

| Metric               | Description   |
|----------------------|---|
| CachePercentDirty    | The overall percentage of the gateway cache that has not been persisted to Amazon. The sample is taken at the end of the reporting period.  Use this metric with the Sum statistic. |
|                      | Ideally, this metric should remain low.   |
|                      | Units: Percent  |
| CachePercentUsed     | The overall percent of the gateway cache storage that is used. The sample is taken at the end of the reporting period.  |
|                      | Units: Percent  |
| CacheUsed            | This metric reports the number of used bytes in the gateway cache.  |
|                      | Units: Bytes  |
| CloudBytesDownloaded | The total number of bytes that the gateway downloaded from Amazon during the reporting period.  |
|                      | Use this metric with the Sum statistic to measure throughput and with the Samples statistic to measure IOPS.  |
|                      | Units: Bytes  |

| Metric                | Description   |
|-----------------------|---|
| CloudBytesUploaded    | The total number of bytes that the gateway uploaded to Amazon during the reporting period.  |
|                       | Use this metric with the Sum statistic to measure throughput and with the Samples statistic to measure input/output operations per second (IOPS).   |
|                       | Units: Bytes  |
| FilesFailingUpload    | This metric tracks the number of files which are failing to upload to Amazon. These files will generate health notifications which contain more information on the issue.   |
|                       | Use this metric with the Sum statistic to show<br>the number of files which are currently failing<br>to upload to Amazon.   |
|                       | Units: Count  |
| FileSharesUnavailable | This metric provides the number of file shares on this gateways which are in the <b>Unavailable</b> state.  |
|                       | 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. |
|                       | Units: Count  |

| Metric              | Description  |
|---------------------|--|
| FilesRenamed        | This metric tracks the number of files renamed in the reporting period.  |
|                     | Units: Count   |
| HealthNotifications | This metric reports the number of health notifications that were generated by this gateway in the reporting period.  |
|                     | Units: Count   |
| IndexEvictions      | This metric reports the number of files whose metadata was evicted from the cached index of file metadata to make room for new entries. The gateway maintains this metadata index, which is populated from the Amazon Cloud on demand. |
|                     | Units: Count   |
| IndexFetches        | This metric reports the number of files for which metadata was fetched. The gateway maintains a cached index of file metadata, which is populated from the Amazon Cloud on demand.   |
|                     | Units: Count   |
| IoWaitPercent       | This metric reports the percentage of time that the CPU is waiting for a response from the local disk.   |
|                     | Units: Percent   |
| MemTotalBytes       | This metric reports the total amount of memory on the gateway.   |
|                     | Units: Bytes   |

| Metric                  | Description  |
|-------------------------|--|
| MemUsedBytes            | This metric reports the amount of used memory on the gateway.  |
|                         | Units: Bytes   |
| NfsSessions             | This metric reports the number of NFS sessions that are active on the gateway.   |
|                         | Units: Count   |
| RootDiskFreeBytes       | This metric reports the number of available bytes on the root disk of the gateway.   |
|                         | If this metric reports less than 20 GB are free, you should increase the size of the root disk.  |
|                         | To increase the root disk size, you can increase<br>the size of existing root disk on the VM. When<br>the VM is rebooted, gateway recognizes the<br>increased size on the root disk. |
|                         | Units: Bytes   |
| S3Get0bjectRequestTime  | This metric reports the time for the gateway to complete S3 get object requests.   |
|                         | Units: Milliseconds  |
| S3Put0bjectRequestTime  | This metric reports the time for the gateway to complete S3 put object requests.   |
|                         | Units: Milliseconds  |
| S3UploadPartRequestTime | This metric reports the time for the gateway to complete S3 upload part requests.  |
|                         | Units: Milliseconds  |

| 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 <u>Using Amazon CloudWatch metrics</u>.



### Note

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

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

| Metric            | Description   |
|-------------------|---|
| CacheHitPercent   | Percent of application read operations from<br>the file shares that are served from cache. The<br>sample is taken at the end of the reporting<br>period.  |
|                   | When there are no application read operation s from the file share, this metric reports 100 percent.  |
|                   | Units: Percent  |
| CachePercentDirty | The file share's contribution to the overall percentage of the gateway's cache that has not been persisted to Amazon. The sample is taken at the end of the reporting period.  Use this metric with the Sum statistic.  Ideally, this metric should remain low. |
|                   | (3) Note  Use the CachePercentDirty metric of the gateway to view the overall percentage of the gateway's cache that has not been persisted to Amazon.  |

| Metric               | Description   |
|----------------------|---|
|                      | Units: Percent  |
| CachePercentUsed     | The file share's contribution to the overall percent use of the gateway's cache storage.  The sample is taken at the end of the reporting period. |
|                      | Use the CachePercentUsed metric of the gateway to view overall percent use of the gateway's cache storage.  |
|                      | Units: Percent  |
| CloudBytesUploaded   | The total number of bytes that the gateway uploaded to Amazon during the reporting period.  |
|                      | Use this metric with the Sum statistic to measure throughput and with the Samples statistic to measure IOPS.                                      |
|                      | Units: Bytes  |
| CloudBytesDownloaded | The total number of bytes that the gateway downloaded from Amazon during the reporting period.  |
|                      | Use this metric with the Sum statistic to measure throughput and with the Samples statistic to measure input/output operations per second (IOPS). |
|                      | Units: Bytes  |
|                      |   |

| Metric             | Description   |
|--------------------|---|
| FilesFailingUpload | This metric tracks the number of files which are failing to upload to Amazon. These files will generate health notifications which contain more information on the issue.  Use this metric with the Sum statistic to show the number of files which are currently failing |
|                    | to upload to Amazon.  |
|                    | Units: Count  |
| ReadBytes          | The total number of bytes read from your on-<br>premises applications in the reporting period<br>for a file share.  |
|                    | Use this metric with the Sum statistic to measure throughput and with the Samples statistic to measure IOPS.  |
|                    | Units: Bytes  |
| WriteBytes         | The total number of bytes written to your on-<br>premises applications in the reporting period.   |
|                    | Use this metric with the Sum statistic to measure throughput and with the Samples statistic to measure IOPS.  |
|                    | Units: Bytes  |

# **Understanding 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.                   |
| clientUid (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<br>data | Write<br>data | Create<br>folder | Create<br>file | Rename<br>file/<br>fold<br>er |   | S | Write<br>attribute<br>s<br>(chown) | S | Write<br>attribute<br>s<br>(chgrp) |
|---------------------------------|--------------|---------------|------------------|----------------|-------------------------------|---|---|------------------------------------|---|------------------------------------|
| access<br>e                     |              |               | X                | X              |                               |   |   |                                    | Х |                                    |
| accoun<br>main<br>(SMB<br>only) | X            | X             | X                | X              | X                             | X | Х | X                                  | X | X                                  |
| accoun<br>me<br>(SMB<br>only)   | X            | X             | X                | X              | X                             | X | X | X                                  | X | X                                  |
| bucket                          | Х            | X             | X                | X              | X                             | X | X | X                                  | Х | X                                  |
| client (NFS only)               | X            | X             | X                | X              | X                             | X |   | X                                  | X | X                                  |
| client (NFS only)               | X            | X             | X                | X              | X                             | Х |   | X                                  | X | X                                  |

|                  | Read<br>data | Write<br>data | Create<br>folder | Create<br>file | Rename<br>file/<br>fold<br>er |   | Write attribute s (change ACL - SMB only) | s | Write<br>eattribute<br>s<br>(chmod) | s |
|------------------|--------------|---------------|------------------|----------------|-------------------------------|---|---|---|-------------------------------------|---|
| ctime            |              |               | Х                | Х              |                               |   |   |   |                                     |   |
| groupI           |              |               | X                | X              |                               |   |   |   |                                     |   |
| fileSi<br>nBytes |              |               |                  | X              |                               |   |   |   |                                     |   |
| gatewa           | Х            | X             | Х                | Х              | Х                             | Χ | X   | X | Х                                   | X |
| mtime            |              |               | X                | X              |                               |   |   |   |                                     |   |
| newObj<br>Name   |              |               |                  |                | X                             |   |   |   |                                     |   |
| object<br>e      | X            | X             | X                | X              | X                             | X | X   | X | X                                   | X |
| object<br>e      | X            | X             | X                | Х              | X                             | X | Х   | X | Х                                   | X |
| operat           | X            | X             | Х                | Х              | X                             | X | X   | Х | Х                                   | X |
| ownerI           |              |               | X                | Х              |                               |   |   | Х |                                     |   |
| securi<br>escrip |              |               |                  |                |                               |   | Χ   | X |                                     |   |
| (SMB<br>only)    |              |               |                  |                |                               |   |   |   |                                     |   |

|                | Read<br>data | Write<br>data | Create<br>folder | Create<br>file | Rename<br>file/<br>fold<br>er |   | attribute<br>s | S | Write<br>attribute<br>s<br>(chmod) | Write<br>attribute<br>s<br>(chgrp) |
|----------------|--------------|---------------|------------------|----------------|-------------------------------|---|----------------|---|------------------------------------|------------------------------------|
| shareN         | X            | X             | X                | X              | X                             | X | X              | X | X                                  | X                                  |
| source         | Х            | Χ             | X                | Χ              | Χ                             | X | Х              | Х | X                                  | X                                  |
| source<br>ress | X            | X             | Х                | X              | X                             | X | X              | X | X                                  | X                                  |
| status         | Х            | Χ             | X                | Χ              | Χ                             | X | Х              | Х | X                                  | Х                                  |
| timest         | X            | X             | X                | X              | Х                             | X | X              | X | X                                  | X                                  |
| versic         | Х            | X             | X                | X              | X                             | X | Х              | Х | Χ                                  | Х                                  |

# Maintaining your gateway

Maintaining your gateway includes tasks such as configuring cache storage and upload buffer space, and doing general maintenance to optimize your gateway's performance. These tasks are common to all gateway types.

### **Topics**

- Shutting down your gateway VM
- Managing local disks for your gateway
- Managing Bandwidth for Your Amazon S3 File Gateway
- Managing Gateway Updates in the Amazon Storage Gateway Console
- Performing maintenance tasks using the local console
- Deleting your gateway and removing associated resources

# 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. When you power on your VM, remember to restart the gateway that runs on the VM.

### 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.

# Managing local disks for your gateway

The gateway virtual machine (VM) uses the local disks that you allocate on-premises for buffering and storage. Gateways created on Amazon EC2 instances use Amazon EBS volumes as local disks.

### **Topics**

- Deciding the amount of local disk storage
- Determining the size of cache storage to allocate
- Configuring additional cache storage
- Using ephemeral storage with EC2 gateways

### Deciding the amount of local disk storage

The number and size of disks that you want to allocate for your gateway is up to you. File Gateways require at least one 150 GiB disk to use as a cache. The cache storage acts as the on-premises durable store for data that is pending upload to Amazon S3 or file system. After the initial configuration and deployment of your gateway, you can add more disks for cache storage as your workload demands increase.



### 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.

# Determining the size of cache storage to allocate

Your gateway uses its cache storage to provide low-latency access to your recently accessed data. The cache storage acts as the on-premises durable store for data that's pending upload to Amazon S3.

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 guotas 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.

### **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

- 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.amazonaws.cn/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**.

Add cache storage API Version 2013-06-30 152 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 for Linux Instances.

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 <u>Getting</u> notified about file operations.
- 3. Call the <u>NotifyWhenUploaded API</u> 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.

# Managing Bandwidth for Your Amazon S3 File Gateway

You can limit the upload throughput from your gateway to Amazon 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 Amazon Web Services Management Console, an Amazon Software Development Kit (SDK), or the Amazon Storage Gateway API (see <a href="UpdateBandwidthRateLimitSchedule">UpdateBandwidthRateLimitSchedule</a> in the Amazon 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 <a href="View and edit the bandwidth-rate-limit schedule for your gateway using the Storage Gateway console">View and edit the bandwidth-rate-limit schedule for your gateway using the Storage Gateway console</a>.

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

Managing Bandwidth API Version 2013-06-30 154

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**

- <u>View and edit the bandwidth-rate-limit schedule for your gateway using the Storage Gateway</u> console
- Updating Gateway Bandwidth-Rate Limits Using the Amazon SDK for Java
- Updating Gateway Bandwidth-Rate Limits Using the Amazon SDK for .NET
- Updating Gateway Bandwidth-Rate Limits Using the Amazon 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 <a href="https://console.amazonaws.cn/storagegateway/home">https://console.amazonaws.cn/storagegateway/home</a>.
- 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.

(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.

When finished, choose **Save changes**.

## **Updating Gateway Bandwidth-Rate Limits Using the Amazon 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 Amazon 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 Amazon SDK for Java Developer Guide.

### **Example: Updating Gateway Bandwidth-Rate Limits Using the Amazon 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 Amazon service endpoints that you can use with Storage Gateway, see Amazon Storage Gateway Endpoints and Quotas in the Amazon Web Services 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 sqClient;
      // 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
           sqClient = new AWSStorageGatewayClient(new PropertiesCredentials(
```

```
UpdateBandwidthExample.class.getResourceAsStream("AwsCredentials.properties")));
           sqClient.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 =
sqClient.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 bandwith.\n" +
ex.toString());
}
}
}
```

# **Updating Gateway Bandwidth-Rate Limits Using the Amazon 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 Amazon 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 <a href="Getting Started">Getting Started</a> in the Amazon SDK for .NET Developer Guide.

### Example: Updating Gateway Bandwidth-Rate Limits by Using the Amazon 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 Amazon service endpoints that you can use with Storage Gateway, see <a href="Amazon Storage Gateway Endpoints">Amazon Storage Gateway Endpoints and Quotas</a> in the Amazon Web Services 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
                sqConfig = new AmazonStorageGatewayConfig();
                sgConfig.ServiceURL = serviceURL;
                sqClient = new AmazonStorageGatewayClient(sqConfig);
                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
updateBandwidthRateLimitScheuduleResponse =
sqClient.UpdateBandwidthRateLimitSchedule(updateBandwidthRateLimitScheduleRequest);
                   String returnGatewayARN =
updateBandwidthRateLimitScheuduleResponse.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 bandwith.\n" +
ex.ToString());
           }
       }
   }
```

# Updating Gateway Bandwidth-Rate Limits Using the Amazon 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 Amazon 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 Amazon Tools for Windows PowerShell User Guide.

# Example: Updating Gateway Bandwidth-Rate Limits by Using the Amazon 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:</pre>
```

```
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/userquide/
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)
```

# Managing Gateway Updates in the Amazon Storage Gateway Console

Storage Gateway periodically releases important software updates for your gateway. You can manually apply updates on the Storage Gateway Management Console, or wait until the updates are automatically applied during the configured maintenance schedule. Although Storage Gateway

Managing Gateway Updates API Version 2013-06-30 162

checks for updates every minute, it only goes through maintenance and restarts if there are updates.

Gateway software releases regularly include operating system updates and security patches that have been validated by Amazon. These updates are typically released every six months, and are applied as part of the normal gateway update process during scheduled maintenance windows.



### Note

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 gateway update mechanism (for example, SSM or hypervisor tools) may cause the gateway to malfunction.

Before any update is applied to your gateway, Amazon notifies you with a message on the Storage Gateway console and your Amazon Health Dashboard. For more information, see Amazon Health Dashboard. The VM doesn't reboot, but the gateway is unavailable for a short period while it's being updated and restarted.

When you deploy and activate your gateway, a default weekly maintenance schedule is set. You can modify the maintenance schedule at any time. When updates are available, the **Details** tab displays a maintenance message. You can see the date and time that the last successful update was applied to your gateway on the **Details** tab.

### To modify the maintenance schedule

- 1. Open the Storage Gateway console at https://console.amazonaws.cn/storagegateway/home.
- On the navigation pane, choose **Gateways**, and choose the gateway that you want to modify 2. the update schedule for.
- For **Actions**, choose **Edit maintenance window** to open the Edit maintenance start time dialog box.
- For **Schedule**, choose **Weekly** or **Monthly** to schedule updates.
- If you choose Weekly, modify the values for Day of the week and Time.

If you choose **Monthly**, modify the values for **Day of the month** and **Time**. If you choose this option and you get an error, it means your gateway is an older version and has not been upgraded to a newer version yet.



### Note

The maximum value that can be set for day of the month is 28. If 28 is selected, the maintenance start time will be on the 28th day of every month.

Your maintenance start time appears on the **Details** tab for the gateway next time that you open the **Details** tab.

## Performing maintenance tasks using the local console

You can perform the following maintenance tasks using the host's local console. Local console tasks can be performed on the VM host or the Amazon EC2 instance. Many of the tasks are common among the different hosts, but there are also some differences.

### **Topics**

- Accessing the Gateway Local Console
- Performing tasks on the virtual machine local console
- Performing tasks on the Amazon EC2 gateway local console
- Configuring Network Adapters for Your Gateway

### **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
```

You can choose available VMs by Id.

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

```
# virsh console VM_Id
```

```
[[root@localhost vms]# virsh console 7
Connected to domain SGW_KVM
Escape character is ^]

Appliance
Login to change your network configuration and other settings.
localhost login: _
```

- 3. To get default credentials to log in to the local console, see <u>Logging in to the File Gateway</u> local console.
- 4. After you have logged in, you can activate and configure your gateway.

```
Appliance Activation - Configuration
## Currently connected network adapters:
##
##
 eth0: 10.0.3.32
1: HTTP/SOCKS Proxy Configuration
2: Network Configuration
3: Test Network Connectivity
4: View System Resource Check (0 Errors)
5: System Time Management
6: License Information
7: Command Prompt
0: Get activation key
Press "x" to exit session
Enter command:
```

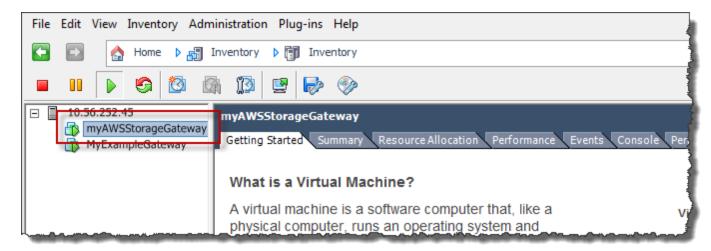
### Accessing the Gateway Local Console with VMware ESXi

### To access your gateway's local console with VMware ESXi

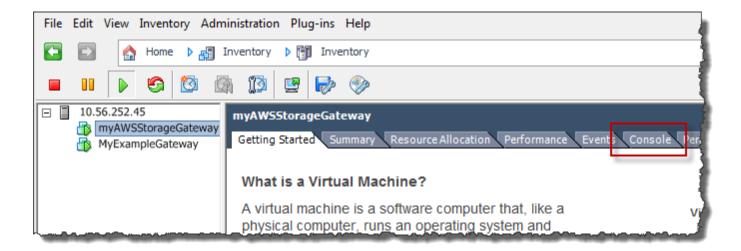
- In the VMware vSphere client, select your gateway VM.
- Make sure that the gateway is turned on.



If your gateway VM is turned on, a green arrow icon appears with the VM icon, as shown in the following screen shot. If your gateway VM is not turned on, you can turn it on by choosing the green **Power On** icon on the **Toolbar** menu.



### Choose the Console tab.



After a few moments, the VM is ready for you to log in.

Note

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



4. To log in using the default credentials, continue to the procedure <u>Logging in to the File</u> Gateway local console.

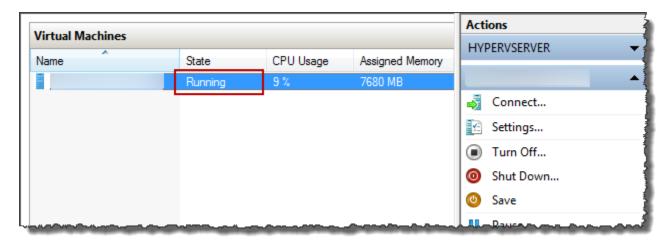
### Access the Gateway Local Console with Microsoft Hyper-V

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

- 1. In the Virtual Machines list of the Microsoft Hyper-V Manager, select your gateway VM.
- 2. Make sure that the gateway is turned on.

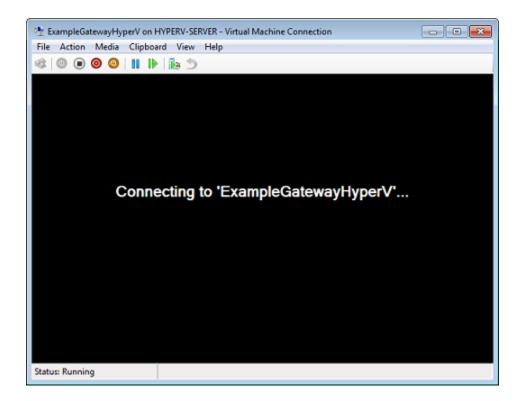
### Note

If your gateway VM is turned on, Running is displayed as the **State** of the VM, as shown in the following screen shot. If your gateway VM is not turned on, you can turn it on by choosing **Start** in the **Actions** pane.



3. In the **Actions** pane, choose **Connect**.

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 VM is ready for you to log in.

```
Login to change your network configuration and other gateway settings.

For more information, please see:
https://docs.aws.amazon.com/console/storagegateway/LocalConsole

localhost login: _
```

4. To log in using the default credentials, continue to the procedure <u>Logging in to the File</u> Gateway 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
- Configuring an HTTP proxy
- Configuring your gateway network settings
- Testing your gateway's network connectivity
- Viewing your gateway system resource status
- Configuring a Network Time Protocol (NTP) server for your gateway
- Running Storage Gateway commands on the local console

# 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 local console, you use the default sign-in credentials to log in. These default login credentials give you access to menus where you can configure gateway network settings and change the password from the local console. Amazon Storage Gateway allows you to set your own password from the Storage Gateway console instead of changing the password from

the local console. You don't need to know the default password to set a new password. For more information, see Setting the local console password from the Storage Gateway console.

#### To log in to the gateway's local console

If this is your first time logging in to the local console, log in to the VM with the default credentials. The default user name is admin and the password is password. Otherwise, use your credentials to log in.



#### Note

We recommend changing the default password by entering the corresponding numeral for Gateway Console from the Amazon 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.

#### Setting the local console password from the Storage Gateway console

When you log in to the local console for the first time, you log in to the VM with the default credentials. For all types of gateways, you use default credentials. The user name is admin and the password is password.

We recommend that you always set a new password immediately after you create your new gateway. You can set this password from the Amazon Storage Gateway console rather than the local console if you want. You don't need to know the default password to set a new password.

### To set the local console password on the Storage Gateway console

- Open the Storage Gateway console at https://console.amazonaws.cn/storagegateway/home. 1.
- On the navigation pane, choose **Gateways**, and then choose the gateway for which you want 2. to set a new password.
- 3. For **Actions**, choose **Set Local Console Password**.
- In the **Set Local Console Password** dialog box, enter a new password, confirm the password, and then choose Save.

Your new password replaces the default password. Storage Gateway doesn't save the password but rather safely transmits it to the VM.



#### 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 Amazon 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

- 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.
- From the Amazon Appliance Activation Configuration main menu, enter the corresponding numeral to select **Configure HTTP Proxy**.

- 3. From the **Amazon 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 <u>Accessing the</u> Gateway Local Console with VMware ESXi.
  - For more information on logging in to the Microsoft Hyper-V local console, see <u>Access the</u> Gateway Local Console with Microsoft Hyper-V.
  - For more information on logging in to the KVM local console, see <u>Accessing the Gateway</u> Local Console with Linux KVM.
- 2. From the **Amazon Appliance Activation Configuration** main menu, enter the corresponding numeral to select **Network Configuration**.
- 3. From the **Network Configuration** menu, perform one of the following tasks:

| To Perform This Task                       | Do This  |
|--|--|
| Get information about your network adapter | Enter the corresponding numeral to select  Describe Adapter. |

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

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

| To Perform This Task                        | Do This  |
|---|--|
|   | For example, suppose that your gateway VM uses two interfaces configured as DHCP. If you later set one interface to a static IP, the other interface is deactivated. To activate the interface in this case, you must set it to a static IP.  If both interfaces are initially set to use static |
|   | IP addresses and you then set the gateway to use DHCP, both interfaces use DHCP.   |
| Configure a hostname for your gateway       | Enter the corresponding numeral to select <b>Configure Hostname</b> .  |
|   | You are prompted to choose whether the gateway will use a static hostname that you specify, or aquire one automatically through DCHP or rDNS.  |
|   | (i) Note  If you configure a static hostname for your gateway, you must create an A record in your DNS system that points the gateway's IP address to its static hostname.   |
| View your gateway's hostname configura tion | Enter the corresponding numeral to select  View Hostname Configuration.  |
|   | Your gateway's hostname, aquisition mode, domain, and Active Directory realm are displayed.  |

| To Perform This Task                                   | Do This  |
|--|--|
| Reset all your gateway's network configuration to DHCP | Enter the corresponding numeral to select  Reset all to DHCP.  All network interfaces are set to use DHCP.   |
|  | ★ Important If your gateway has already been activated, you must shut down and restart your gateway from the Storage Gateway console for the settings to take effect. For more information, see Shutting down your gateway VM. |
| Set your gateway's default route adapter               | Enter the corresponding numeral to select <b>Set Default Adapter</b> .  The available adapters for your gateway are shown, and you are prompted to choose one of the adapters—for example, <b>eth0</b> .                       |
| Edit your gateway's DNS configuration                  | Enter the corresponding numeral to select  Edit DNS Configuration.  The available adapters of the primary and secondary DNS servers are displayed. You are prompted to provide the new IP address.                             |

| To Perform This Task                  | Do This  |
|---------------------------------------|--|
| View your gateway's DNS configuration | Enter the corresponding numeral to select  View DNS Configuration.  The available adapters of the primary and secondary DNS servers are displayed.  (i) Note  For some versions of the VMware hypervisor, you can edit the adapter configuration in this menu. |
| View routing tables                   | Enter the corresponding numeral to select  View Routes.  The default route of your gateway is displayed.   |

# Testing your gateway's network connectivity

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

# To test your gateway's network connectivity

- 1. Log in to your gateway's local console:
  - For more information on logging in to the VMware ESXi local console, see <u>Accessing the Gateway Local Console with VMware ESXi.</u>
  - For more information on logging in to the Microsoft Hyper-V local console, see <u>Access the</u> Gateway Local Console with Microsoft Hyper-V.
  - For more information on logging in to the KVM local console, see <u>Accessing the Gateway</u> Local Console with Linux KVM.

- 2. From the **Amazon 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 Amazon Web Services 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 Amazon Web Services Region that you want to test. For supported Amazon Web Services Regions and a list of Amazon service endpoints you can use with Storage Gateway, see <a href="#">Amazon</a> Storage Gateway endpoints and quotas in the Amazon Web Services 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 <u>Accessing the Gateway</u> Local Console with VMware ESXi.
  - For more information on logging in to the Microsoft Hyper-V local console, see <u>Access the</u> Gateway Local Console with Microsoft Hyper-V.

- For more information on logging in to the KVM local console, see <u>Accessing the Gateway</u> Local Console with Linux KVM.
- 2. From the **Amazon 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

Log in to your gateway's local console:

- For more information on logging in to the VMware ESXi local console, see <u>Accessing the</u> Gateway Local Console with VMware ESXi.
- For more information on logging in to the Microsoft Hyper-V local console, see <u>Access the</u> Gateway Local Console with Microsoft Hyper-V.
- For more information on logging in to the KVM local console, see <u>Accessing the Gateway</u> Local Console with Linux KVM.
- 2. From the **Amazon Appliance Activation Configuration** main menu, enter the corresponding numeral to select **System Time Management**.
- 3. From the **System Time Management** menu, enter the corresponding numeral to perform one of the following tasks.

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

| To Perform This Task               | Do This  |
|------------------------------------|--|
|                                    | 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 <a href="Synchronizing VM Time with Host Time">Synchronizing VM Time with Host Time</a> . |
|                                    | For a gateway deployed on Microsoft Hyper-V, you should periodically check your VM's time. For more information, see <a href="Synchronizing Your Gateway VM Time">Synchronizing Your Gateway VM Time</a> .   |
|                                    | For a gateway deployed on KVM, you can check and synchronize the VM time using virsh command line interface for KVM.   |
| Edit your NTP server configuration | Enter the corresponding numeral to select  Edit NTP Configuration.  You are prompted to provide a preferred and  |
|                                    | a secondary NTP server.  |
| View your NTP server configuration | Enter the corresponding numeral to select <b>View NTP Configuration</b> .  |
|                                    | Your NTP server configuration is displayed.  |

# **Running Storage Gateway commands on the local console**

The VM local console in Storage Gateway helps provide a secure environment for configuring and diagnosing issues with your gateway. Using the local console commands, you can perform maintenance tasks such as saving routing tables, connecting to Amazon Web Services 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 <u>Accessing the</u> Gateway Local Console with VMware ESXi.
  - For more information on logging in to the Microsoft Hyper-V local console, see <u>Access the</u> Gateway Local Console with Microsoft Hyper-V.
  - For more information on logging in to the KVM local console, see <u>Accessing the Gateway</u> Local Console with Linux KVM.
- 2. From the **Amazon Appliance Activation Configuration** main menu, enter the corresponding numeral to select **Gateway Console**.
- 3. From the gateway console command prompt, enter **h**.

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

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

| Command              | Function  |
|----------------------|---|
| ip                   | Show / manipulate routing, devices, and tunnels.  |
|                      | We recommend configuring network or IP settings using the Storage Gateway console or the dedicated local console menu option. For instructions, see <a href="Configuring your gateway network">Configuring your gateway network</a> 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 t roubleshooting.   |
| open-support-channel | Connect to Amazon Support. For instructi ons on how to turn on Amazon support access, see You want Amazon 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.   |
| tcptraceroute        | Collect traceroute output on TCP traffic to a destination.  |

| Command  | Function  |  |
|----------|---|--|
| sslcheck | Returns output with certificate issuer  |  |
|          | Storage Gateway uses certificate issuer verification for its working and doesn't support ssl inspection. If this command has an issuer other than aws-appliance@amazon.com, then it is likely that there is an appliance performing an ssl inspection. In that case, you can bypass ssl inspection for Storage Gateway Appliance. |  |

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

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

# Performing tasks on the Amazon EC2 gateway local console

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

#### **Topics**

- Logging in to your Amazon EC2 gateway local console
- Routing your gateway deployed on Amazon EC2 through an HTTP proxy
- Testing your gateway's network connectivity
- Viewing your gateway system resource status
- Running Storage Gateway commands on the local console for an Amazon EC2 gateway
- Configuring your Amazon EC2 gateway network settings

# Logging in to your Amazon EC2 gateway local console

You can connect to your Amazon EC2 instance by using a Secure Shell (SSH) client. For detailed information, see <u>Connect to your instance</u> 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 <u>Amazon EC2 key pairs</u> 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 **Amazon 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                | <u>Viewing your gateway system resource statu</u> <u>s</u> .                    |
| 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 Amazon.

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 Amazon 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 <u>Logging in to your Amazon EC2</u> gateway local console.
- From the Amazon Appliance Activation Configuration main menu, enter the corresponding numeral to select Configure HTTP Proxy.
- 3. From the **Amazon 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 <u>Logging in to your Amazon EC2</u> <u>gateway local console</u>.
- 2. From the **Amazon 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 Amazon Web Services 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 Amazon Web Services Region that you want to test. For supported Amazon Web Services Regions and a list of Amazon service endpoints you can use with Storage Gateway, see <a href="#">Amazon Storage Gateway endpoints and quotas</a> in the Amazon Web Services 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 instructions, see <u>Logging in to your Amazon EC2</u> gateway local console.
- 2. From the **Amazon 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.

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

The Amazon 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 Amazon Web Services Support.

#### To run a configuration or diagnostic command

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

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

| Command  | Function   |  |
|----------|--|--|
| dig      | Collect output from dig for DNS troublesh ooting.  |  |
| exit     | Return to Configuration menu.  |  |
| h        | Display available command list.  |  |
| ifconfig | View or configure network interfaces.  Note  We recommend configuring network or IP settings using the Storage Gateway console or the dedicated local console menu option. For instructions, see Configuring your gateway network settings.            |  |
| ip       | Show / manipulate routing, devices, and tunnels.  Note  We recommend configuring network or IP settings using the Storage Gateway console or the dedicated local console menu option. For instructions, see Configuring your gateway network settings. |  |
| iptables | Administration tool for IPv4 packet filtering and NAT.   |  |

| Command              | Function   |
|----------------------|--|
| ncport               | Test connectivity to a specific TCP port on a network.     |
| nping                | Collect output from nping for network t roubleshooting.    |
| open-support-channel | Connect to Amazon Support.                                 |
| save-iptables        | Persist IP tables.   |
| save-routing-table   | Save newly added routing table entry.                      |
| tcptraceroute        | Collect traceroute output on TCP traffic to a destination. |

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

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

# **Configuring your Amazon EC2 gateway network settings**

You can view and configure your gateway's network settings through the local console.

# To configure your network settings

- 1. Log in to your gateway's local console. For instructions, see <u>Logging in to your Amazon EC2</u> gateway local console.
- From the Amazon Appliance Activation Configuration main menu, enter the corresponding numeral to select Network Configuration.
- 3. From the **Amazon Appliance Activation Network Configuration** menu, enter the corresponding numeral for the task you want to perform:
  - Edit DNS Configuration The available adapters of the primary and secondary DNS servers are displayed. You are prompted to provide the new IP address.
  - View DNS Configuration The available adapters of the primary and secondary DNS servers are displayed.

• Configure Hostname - You are prompted to choose whether the gateway will use a static hostname that you specify, or aquire one automatically through DCHP or rDNS.



#### Note

If you configure a static hostname for your gateway, you must create an A record in your DNS system that points the gateway's IP address to its static hostname.

• View Hostname Configuration - Your gateway's hostname, aguisition mode, domain, and Active Directory realm are displayed.

# **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 Amazon.

In a typical multiple-adapter use case, one adapter is configured as the route by which the gateway communicates with Amazon (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 Amazon, then file share traffic for that target and Amazon 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.
  - 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.

- In the Storage Gateway console, turn on the gateway. 7.
- In the **Navigation** pane of the Storage Gateway console, choose **Gateways** and choose the 8. 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.

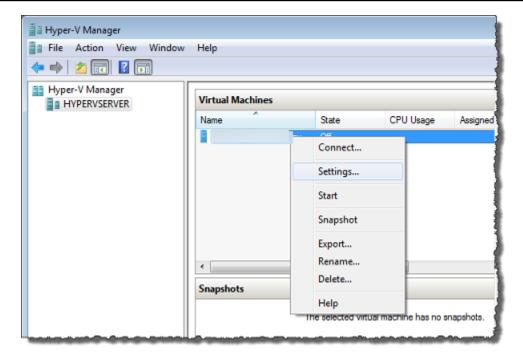
For information about local console tasks common to VMware, Hyper-V, and KVM hosts, see Performing tasks on the virtual machine local console

# Configuring Your Gateway for Multiple NICs in Microsoft Hyper-V Host

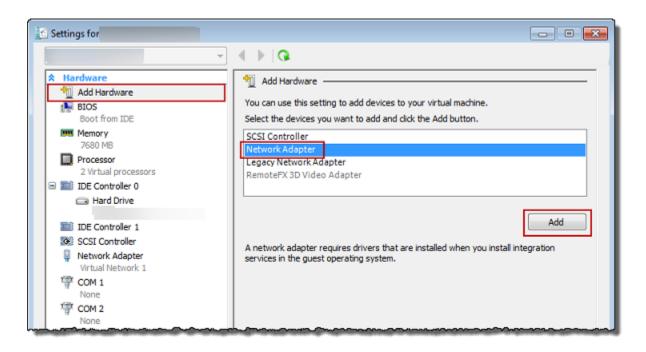
The following procedure assumes that your gateway VM already has one network adapter defined and that you are adding a second adapter. This procedure shows how to add an adapter for a Microsoft Hyper-V host.

# To configure your gateway to use an additional network adapter in a Microsoft Hyper-V Host

- On the Storage Gateway console, turn off the gateway. 1.
- In the Microsoft Hyper-V Manager, select your gateway VM. 2.
- If the VM isn't turned off already, open the context (right-click) menu for your gateway and choose Turn Off.
- In the client, open the context menu for your gateway VM and choose **Settings**.

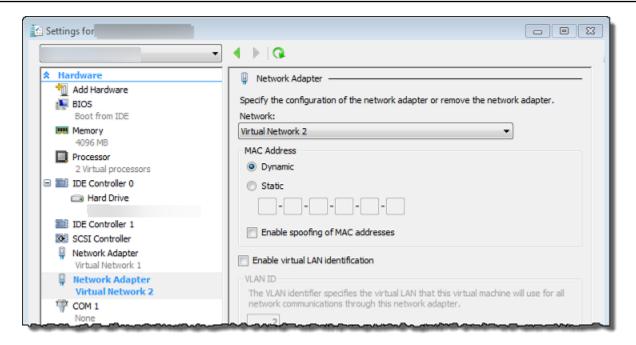


- 5. In the **Settings** dialog box for the VM, for **Hardware**, choose **Add Hardware**.
- 6. In the **Add Hardware** pane, choose **Network Adapter**, and then choose **Add** to add a device.



7. Configure the network adapter, and then choose **Apply** to apply settings.

In the following example, **Virtual Network 2** is selected for the new adapter.



- 8. In the **Settings** dialog box, for **Hardware**, confirm that the second adapter was added, and then choose **OK**.
- 9. On the Storage Gateway console, turn on the gateway.
- 10. In the **Navigation** pane choose **Gateways**, then select 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

# 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 Amazon 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 *Amazon Storage Gateway API Reference*.

#### **Topics**

- Deleting Your Gateway by Using the Storage Gateway Console
- Removing Resources from a Gateway Deployed On-Premises
- Removing Resources from a Gateway Deployed on an Amazon EC2 Instance

# **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.



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.amazonaws.cn/storagegateway/home.
- 2. Choose **Gateways**, then select one or more gateways to delete.
- 3. For **Actions**, choose **Delete gateway**. The confirmation dialog box appears.

#### Marning

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.

- Verify that you want to delete the specified gateways, then type the word *delete* in the confirmation box, and choose **Delete**.
- (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.

# Removing Resources from a Gateway Deployed On-Premises

You can use the instructions following to remove resources from a gateway that is deployed onpremises.

# Removing Resources from a Volume Gateway Deployed on a VM

If the gateway you want to delete are deployed on a virtual machine (VM), we suggest that you take the following actions to clean up resources:

Delete the gateway.

# Removing Resources from a Gateway Deployed on an Amazon EC2 Instance

If you want to delete a gateway that you deployed on an Amazon EC2 instance, we recommend that you clean up the Amazon resources that were used with the gateway. Doing so helps avoid unintended usage charges.

# Removing Resources from Your Cached Volumes Deployed on Amazon EC2

If you deployed a gateway with cached volumes on EC2, we suggest that you take the following actions to delete your gateway and clean up its resources:

- 1. In the Storage Gateway console, delete the gateway as shown in <u>Deleting Your Gateway by</u> Using the Storage Gateway Console.
- 2. In the Amazon EC2 console, stop your EC2 instance if you plan on using the instance again. Otherwise, terminate the instance. If you plan on deleting volumes, make note of the block devices that are attached to the instance and the devices' identifiers before terminating the instance. You will need these to identify the volumes you want to delete.
- 3. In the Amazon EC2 console, remove all Amazon EBS volumes that are attached to the instance if you don't plan on using them again. For more information, see <u>Clean Up Your Instance and Volume</u> in the *Amazon EC2 User Guide for Linux Instances*.

# Replacing your existing File Gateway with a new instance

You can replace an existing File Gateway with a new instance as your data and performance needs grow, or if you receive an Amazon 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 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-ca che read operations. | Data in cache is downloade d from the Amazon cloud. This method is optimal for write-heavy workloads, if your applications can tolerate the delay caused by out-of-ca che 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  |
|------------|---|--|
|            |   | Writing to one Amazon S3 bucket from two file shares simultaneously is not supported, so all clients must be remapped from one share to the other simultaneously, rather than gradually. |
| 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.

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

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

- Stop any applications that are writing to the existing File Gateway. 1.
- Verify that the CachePercentDirty metric on the Monitoring tab for the existing File Gateway is 0.

Shut down the existing File Gateway by powering off the host virtual machine (VM) using its 3. 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, cache disks, and upload buffer 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 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.

Create a new Amazon 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.

- Start the new Storage Gateway VM.
- 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.

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=sqw-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 Amazon 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 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.



#### Marning

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 File Gateway instance with empty cache disk and new Gateway ID:

- Stop any applications that are writing to the existing File Gateway. Verify that the CachePercentDirty metric on the **Monitoring** tab is 0 before you set up file shares on the new gateway.
- Use the Amazon Command Line Interface (Amazon CLI) to gather and save the configuration information about your existing File Gateway and file shares by doing the following:
  - Save the gateway configuration information for the 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 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 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 NFS 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 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 File Gateway, and the path used by the NFS client to identify the mount point.

3. Create a new File Gateway with the same settings and configuration as the old gateway. If necessary, refer to the information you saved in Step 2.

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

#### Marning

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

Delete the old gateway virtual machine or EC2 instance. 7.

### 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
- Optimizing Gateway Performance
- Using VMware vSphere High Availability with Storage Gateway

### 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.



We don't recommend using ephemeral storage. For information about using ephemeral storage, see <u>Using ephemeral storage with EC2 gateways</u>.

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 <a href="Modifying an EBS volume using elastic volumes">Modifying an EBS volume using elastic volumes</a> (console).

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<br>Configurations                                     | Protocol             | Write<br>throughput (file<br>sizes 1 GB) | Cache hit read throughput | Cache miss read throughput |
|---|----------------------|--|---------------------------|----------------------------|
| Root disk: 80 GB, io1 SSD, 4,000                              | NFSv3 - 1 thread     | 110 MiB/sec<br>(0.92 Gbps)               | 590 MiB/sec (4.9<br>Gbps) | 310 MiB/sec (2.6<br>Gbps)  |
| Cache disk: 512   | NFSv3 - 8<br>threads | 160 MiB/sec (1.3<br>Gbps)                | 590 MiB/sec (4.9<br>Gbps) | 335 MiB/sec (2.8<br>Gbps)  |
| GiB cache, io1,<br>1,500 provision<br>ed IOPS                 | NFSv4 - 1 thread     | 130 MiB/sec (1.1<br>Gbps)                | 590 MiB/sec (4.9<br>Gbps) | 295 MiB/sec (2.5<br>Gbps)  |
| Minimum<br>network  | NFSv4 - 8<br>threads | 160 MiB/sec (1.3<br>Gbps)                | 590 MiB/sec (4.9 Gbps)    | 335 MiB/sec (2.8 Gbps)     |
| performance: 10<br>Gbps                                       | SMBV3 - 1<br>thread  | 115 MiB/sec (1.0<br>Gbps)                | 325 MiB/sec (2.7<br>Gbps) | 255 MiB/sec (2.1<br>Gbps)  |
| CPU: 16 vCPU   RAM: 32 GB  NFS protocol recommended for Linux | SMBV3 - 8<br>threads | 190 MiB/sec (1.6<br>Gbps)                | 590 MiB/sec (4.9<br>Gbps) | 335 MiB/sec (2.8 Gbps)     |
| Storage<br>Gateway  | NFSv3 - 1 thread     | 265 MiB/sec (2.2<br>Gbps)                | 590 MiB/sec (4.9<br>Gbps) | 310 MiB/sec (2.6<br>Gbps)  |
| Hardware<br>Appliance   | NFSv3 - 8<br>threads | 385 MiB/sec (3.1<br>Gbps)                | 590 MiB/sec (4.9 Gbps)    | 335 MiB/sec (2.8 Gbps)     |
| Minimum<br>network<br>performance: 10                         | NFSv4 - 1 thread     | 310 MiB/sec (2.6<br>Gbps)                | 590 MiB/sec (4.9<br>Gbps) | 295 MiB/sec (2.5<br>Gbps)  |
| Gbps  | NFSv4 - 8<br>threads | 385 MiB/sec (3.1<br>Gbps)                | 590 MiB/sec (4.9 Gbps)    | 335 MiB/sec (2.8 Gbps)     |

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

## File Gateway performance on Windows clients

| Example<br>Configurations                     | Protocol             | Write<br>throughput (file<br>sizes 1 GB) | Cache hit read throughput | Cache miss read<br>throughput |
|---|----------------------|--|---------------------------|-------------------------------|
| Root disk: 80 GB, io1 SSD, 4,000              | SMBV3 - 1<br>thread  | 150 MiB/sec (1.3<br>Gbps)                | 180 MiB/sec (1.5<br>Gbps) | 20 MiB/sec (0.2<br>Gbps)      |
| Cache disk: 512                               | SMBV3 - 8<br>threads | 190 MiB/sec (1.6<br>Gbps)                | 335 MiB/sec (2.8<br>Gbps) | 195 MiB/sec (1.6<br>Gbps)     |
| GiB cache, io1,<br>1,500 provisioned<br>IOPS  | NFSv3 - 1<br>thread  | 95 MiB/sec (0.8<br>Gbps)                 | 130 MiB/sec (1.1<br>Gbps) | 20 MiB/sec (0.2<br>Gbps)      |
| Minimum<br>network<br>performance: 10<br>Gbps | NFSv3 - 8<br>threads | 190 MiB/sec (1.6<br>Gbps)                | 330 MiB/sec (2.8 Gbps)    | 190 MiB/sec (1.6<br>Gbps)     |
| CPU: 16 vCPU  <br>RAM: 32 GB                  |                      |  |                           |                               |
| SMB protocol recommended for Windows          |                      |  |                           |                               |
| Storage Gateway Hardware                      | SMBV3 - 1<br>thread  | 230 MiB/sec (1.9<br>Gbps)                | 255 MiB/sec (2.1<br>Gbps) | 20 MiB/sec (0.2<br>Gbps)      |
| Appliance<br>Minimum                          | SMBV3 - 8<br>threads | 835 MiB/sec (7.0<br>Gbps)                | 475 MiB/sec (4.0<br>Gbps) | 195 MiB/sec (1.6<br>Gbps)     |
| network<br>performance: 10<br>Gbps            | NFSv3 - 1<br>thread  | 135 MiB/sec (1.1<br>Gbps)                | 185 MiB/sec (1.6<br>Gbps) | 20 MiB/sec (0.2<br>Gbps)      |
|   | NFSv3 - 8<br>threads | 545 MiB/sec (4.6<br>Gbps)                | 470 MiB/sec (4.0 Gbps)    | 190 MiB/sec (1.6<br>Gbps)     |

| Example<br>Configurations                     | Protocol                      | Write<br>throughput (file<br>sizes 1 GB) | Cache hit read<br>throughput | Cache miss read throughput |
|---|-------------------------------|--|------------------------------|----------------------------|
| Root disk: 80 GB, io1 SSD, 4,000              | SMBV3 - 1<br>thread           | 230 MiB/sec (1.9<br>Gbps)                | 265 MiB/sec (2.2<br>Gbps)    | 30 MiB/sec (0.3<br>Gbps)   |
| Cache disk: 4 x 2                             | SMBV3 - 8<br>threads          | 835 MiB/sec (7.0<br>Gbps)                | 780 MiB/sec (6.5<br>Gbps)    | 250 MiB/sec (2.1<br>Gbps)  |
| disks   | thread (1.1. Gbps) Gbps) Gbps | 30 MiB/sec (0.3<br>Gbps)                 |                              |                            |
| Minimum<br>network<br>performance: 10<br>Gbps | NFSv3 - 8<br>threads          | 545 MiB/sec (4.6<br>Gbps)                | 570 MiB/sec (4.8<br>Gbps)    | 240 MiB/sec (2.0<br>Gbps)  |
| CPU: 32 vCPU  <br>RAM: 244 GB                 |                               |  |                              |                            |
| SMB protocol recommended for Windows          |                               |  |                              |                            |

#### 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<br>Shares<br>Per<br>Gateway | Recommended Gateway Capacity Setting | Recommer<br>ed vCPU<br>Cores                      | Recommer<br>ed RAM | Recommer<br>ed Disk<br>Size |
|----------------------------------|--------------------------------------|---|--------------------|-----------------------------|
| 1-10                             | Small                                | 4 (EC2 instance type <b>m4.xlarge</b> or greater) | 16 GiB             | 80 GiB                      |
| 10-20                            | Medium                               | 8 (EC2 instance type m4.2xlarg e or greater)      | 32 GiB             | 160 GiB                     |
| 20+                              | Large                                | 16 (EC2 instance                                  | 64 GiB             | 240 GiB                     |

| File<br>Shares<br>Per<br>Gateway | Recommened Gateway Capacity Setting | Recommer<br>ed vCPU<br>Cores          | Recommer<br>ed RAM | Recommer<br>ed Disk<br>Size |  |
|----------------------------------|-------------------------------------|---------------------------------------|--------------------|-----------------------------|--|
|                                  |                                     | type<br>m4.4xlarg<br>e or<br>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.

### **Optimizing Gateway Performance**

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

### **Add Resources to Your Gateway**

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

#### Use higher-performance disks

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

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



#### Note

CloudWatch metrics are not available for all gateways. For information about gateway metrics, see Monitoring your S3 File Gateway.

#### Add CPU resources to your gateway host

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

When you add additional CPUs to your gateway host server, you increase the processing capability of the gateway. Doing this allows your gateway to deal with, in parallel, both storing data from your application to your local storage and uploading this data to Amazon S3. Additional CPUs also help ensure that your gateway gets enough CPU resources when the host is shared with other VMs. Providing enough CPU resources has the general effect of improving throughput.

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

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

#### Back gateway virtual disks with separate physical disks

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

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

### **Add Resources to Your Application Environment**

#### Increase the bandwidth between your application server and your gateway

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

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

#### Add CPU resources to your application environment

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

### 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 Amazon automatically recovers from most service interruptions. It generally does this in under 60 seconds with no data loss.

To use VMware HA with Storage Gateway, take the steps listed following.

#### **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**

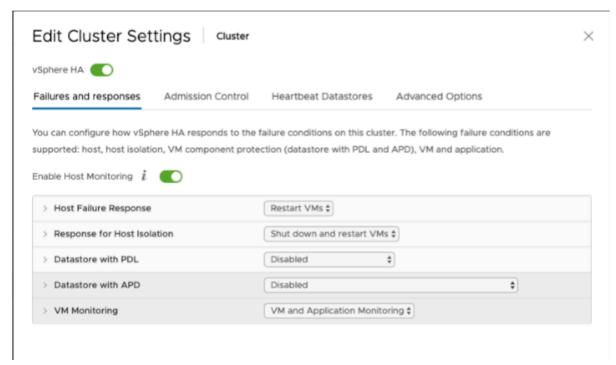
First, if you haven't already created a VMware cluster, create one. For information about how to create a VMware cluster, see Create a vSphere HA Cluster in the VMware documentation.

Next, configure your VMware cluster to work with Storage Gateway.

#### To configure your VMware cluster

- 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 options as listed:
  - · 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

For an example, see the following screen shot.



- 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 <u>Create and activate an Amazon S3 File Gateway</u>

### **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.

Set Up Your Gateway Type API Version 2013-06-30 218

#### 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.

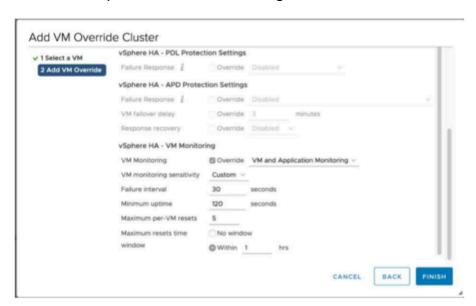
### (Optional) Add Override Options for Other VMs on Your Cluster

If you have other VMs running on your cluster, you might want to set the cluster values specifically for each VM. For instructions, see <u>Customize an Individual Virtual Machine</u> 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**.
- Choose the Configuration tab, and then choose VM Overrides.
- 3. Add a new VM override option to change each value.

For override options, see the following screen shot.



### **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

- Open the Storage Gateway console at https://console.amazonaws.cn/storagegateway/home. 1.
- 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**.
- In the Verify VMware High Availability Configuration box that appears, choose OK. 4.



#### 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.

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.

**Activate Your Gateway** API Version 2013-06-30 220

### **Security in Amazon Storage Gateway**

Cloud security at Amazon is the highest priority. As an Amazon 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 Amazon and you. The <u>shared responsibility model</u> describes this as security *of* the cloud and security *in* the cloud:

- Security of the cloud Amazon is responsible for protecting the infrastructure that runs
   Amazon services in the Amazon Cloud. Amazon also provides you with services that you can use
   securely. Third-party auditors regularly test and verify the effectiveness of our security as part
   of the <u>Amazon Compliance Programs</u>. To learn about the compliance programs that apply to
   Amazon Storage Gateway, see <u>Amazon Services in Scope by Compliance Program</u>.
- **Security in the cloud** Your responsibility is determined by the Amazon service that you use. You are also responsible for other factors including the sensitivity of your data, your company's requirements, and applicable laws and regulations.

This documentation helps you understand how to apply the shared responsibility model when using 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 Amazon services that help you to monitor and secure your Storage Gateway resources.

#### **Topics**

- Data protection in Amazon Storage Gateway
- Identity and access management for Amazon Storage Gateway
- Logging and monitoring in Amazon Storage Gateway
- Compliance validation for Amazon Storage Gateway
- Resilience in Amazon Storage Gateway
- Infrastructure security in Amazon Storage Gateway
- Amazon Security Best Practices

### **Data protection in Amazon Storage Gateway**

The Amazon <u>shared responsibility model</u> applies to data protection in Amazon Storage Gateway. As described in this model, Amazon is responsible for protecting the global infrastructure that runs all of the Amazon Web Services Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. You are also responsible for the security configuration and management tasks for the Amazon Web Services that you use. For more information about data privacy, see the Data Privacy FAQ.

For data protection purposes, we recommend that you protect Amazon Web Services account credentials and set up individual users with Amazon IAM Identity Center or Amazon Identity and Access Management (IAM). That way, each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.
- Use SSL/TLS to communicate with Amazon resources. We require TLS 1.2 and recommend TLS 1.3.
- Set up API and user activity logging with Amazon CloudTrail.
- Use Amazon encryption solutions, along with all default security controls within Amazon Web Services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing sensitive data that is stored in Amazon S3.
- If you require FIPS 140-2 validated cryptographic modules when accessing Amazon through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see Federal Information Processing Standard (FIPS) 140-2.

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form text fields such as a **Name** field. This includes when you work with Storage Gateway or other Amazon Web Services using the console, API, Amazon CLI, or Amazon SDKs. Any data that you enter into tags or free-form text fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

Data protection API Version 2013-06-30 222

### **Data encryption using Amazon KMS**

Storage Gateway uses SSL/TLS (Secure Socket Layers/Transport Layer Security) to encrypt data that is transferred between your gateway appliance and Amazon 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 Amazon Key Management Service (SSE-KMS) keys.

#### **Encrypting a file share**

For a file share, you can configure your gateway to encrypt your objects with Amazon KMS—managed keys by using SSE-KMS. For information on using the Storage Gateway API to encrypt data written to a file share, see CreateNFSFileShare in the Amazon Storage Gateway API Reference.

#### When using Amazon 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 Amazon KMS API operations. For more information, see <u>Using IAM policies with Amazon KMS</u> in the *Amazon Key Management Service Developer Guide*.

### ▲ Important

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

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

### Identity and access management for Amazon Storage Gateway

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

Data encryption API Version 2013-06-30 223

#### **Topics**

- Audience
- Authenticating with identities
- Managing access using policies
- How Amazon Storage Gateway works with IAM
- Identity-based policy examples for Amazon Storage Gateway
- Troubleshooting Amazon 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 Amazon Identity and Access Management (IAM) differs, depending on the work that you do in Amazon SGW.

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

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

### **Authenticating with identities**

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

Audience API Version 2013-06-30 224

If you access Amazon programmatically, Amazon provides a software development kit (SDK) and a command line interface (CLI) to cryptographically sign your requests by using your credentials. If you don't use Amazon tools, you must sign requests yourself. For more information about using the recommended method to sign requests yourself, see <u>Signing Amazon API requests</u> in the *IAM User Guide*.

Regardless of the authentication method that you use, you might be required to provide additional security information. For example, Amazon recommends that you use multi-factor authentication (MFA) to increase the security of your account. To learn more, see <u>Using multi-factor authentication</u> (MFA) in Amazon in the *IAM User Guide*.

#### Amazon Web Services account root user

When you create an Amazon Web Services account, you begin with one sign-in identity that has complete access to all Amazon Web Services and resources in the account. This identity is called the Amazon Web Services account *root user* and is accessed by signing in with the email address and password that you used to create the account. We strongly recommend that you don't use the root user for your everyday tasks. Safeguard your root user credentials and use them to perform the tasks that only the root user can perform. For the complete list of tasks that require you to sign in as the root user, see <u>Tasks that require root user credentials</u> in the *IAM User Guide*.

### **Federated identity**

As a best practice, require human users, including users that require administrator access, to use federation with an identity provider to access Amazon Web Services by using temporary credentials.

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

### IAM users and groups

An <u>IAM user</u> is an identity within your Amazon Web Services account that has specific permissions for a single person or application. Where possible, we recommend relying on temporary credentials instead of creating IAM users who have long-term credentials such as passwords and access keys. However, if you have specific use cases that require long-term credentials with IAM users, we recommend that you rotate access keys. For more information, see <u>Rotate access keys regularly for use cases that require long-term credentials in the IAM User Guide</u>.

An <u>IAM group</u> 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 <a href="When to create an IAM user (instead of a role">When to create an IAM user (instead of a role)</a> in the IAM User Guide.

#### IAM roles

An <u>IAM role</u> is an identity within your Amazon Web Services account that has specific permissions. It is similar to an IAM user, but is not associated with a specific person. You can temporarily assume an IAM role in the Amazon Web Services Management Console by <u>switching roles</u>. You can assume a role by calling an Amazon CLI or Amazon API operation or by using a custom URL. For more information about methods for using roles, see <u>Using IAM roles</u> in the <u>IAM User Guide</u>.

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 <a href="Creating a role for a third-party Identity Provider">Creating a role for a third-party Identity Provider</a> in
  the IAM User Guide.
- **Temporary IAM user permissions** An IAM user or role can assume an IAM role to temporarily take on different permissions for a specific task.
- Cross-account access You can use an IAM role to allow someone (a trusted principal) in a
  different account to access resources in your account. Roles are the primary way to grant crossaccount access. However, with some Amazon Web Services, you can attach a policy directly to a
  resource (instead of using a role as a proxy). To learn the difference between roles and resourcebased policies for cross-account access, see <a href="How IAM roles differ from resource-based policies">How IAM roles differ from resource-based policies</a> in
  the IAM User Guide.
- Cross-service access Some Amazon Web Services use features in other Amazon Web Services. For example, when you make a call in a service, it's common for that service to run applications in Amazon EC2 or store objects in Amazon S3. A service might do this using the calling principal's permissions, using a service role, or using a service-linked role.

- Forward access sessions (FAS) When you use an IAM user or role to perform actions in Amazon, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an Amazon Web Service, combined with the requesting Amazon Web Service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other Amazon Web Services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see <a href="Forward access sessions">Forward access sessions</a>.
- Service role A service role is an <u>IAM role</u> 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 <u>Creating a role to delegate permissions to an Amazon Web Service</u> in the *IAM User Guide*.
- Service-linked role A service-linked role is a type of service role that is linked to an Amazon Web Service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your Amazon Web Services account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.
- Applications running on Amazon EC2 You can use an IAM role to manage temporary credentials for applications that are running on an EC2 instance and making Amazon CLI or Amazon API requests. This is preferable to storing access keys within the EC2 instance. To assign an Amazon role to an EC2 instance and make it available to all of its applications, you create an instance profile that is attached to the instance. An instance profile contains the role and enables programs that are running on the EC2 instance to get temporary credentials. For more information, see <a href="Using an IAM role to grant permissions to applications running on Amazon EC2">Using an IAM role to grant permissions to applications running on Amazon EC2 instances in the IAM User Guide.</a>

To learn whether to use IAM roles or IAM users, see When to create an IAM role (instead of a user) in the IAM User Guide.

### Managing access using policies

You control access in Amazon by creating policies and attaching them to Amazon identities or resources. A policy is an object in Amazon that, when associated with an identity or resource, defines their permissions. Amazon evaluates these policies when a principal (user, root user, or role session) makes a request. Permissions in the policies determine whether the request is allowed or denied. Most policies are stored in Amazon as JSON documents. For more information about the

structure and contents of JSON policy documents, see <u>Overview of JSON policies</u> in the *IAM User Guide*.

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

By default, users and roles have no permissions. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

IAM policies define permissions for an action regardless of the method that you use to perform the operation. For example, suppose that you have a policy that allows the iam: GetRole action. A user with that policy can get role information from the Amazon Web Services Management Console, the Amazon CLI, or the Amazon API.

### **Identity-based policies**

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see <a href="Creating IAM policies">Creating IAM policies</a> in the IAM User Guide.

Identity-based policies can be further categorized as *inline policies* or *managed policies*. Inline policies are embedded directly into a single user, group, or role. Managed policies are standalone policies that you can attach to multiple users, groups, and roles in your Amazon Web Services account. Managed policies include Amazon managed policies and customer managed policies. To learn how to choose between a managed policy or an inline policy, see <a href="Choosing between managed policies and inline policies">Choosing between managed policies and inline policies</a> 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 <u>specify a principal</u> in a resource-based policy. Principals can include accounts, users, roles, federated users, or Amazon Web Services.

Resource-based policies are inline policies that are located in that service. You can't use Amazon managed policies from IAM in a resource-based policy.

### **Access control lists (ACLs)**

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Amazon S3, Amazon WAF, and Amazon VPC are examples of services that support ACLs. To learn more about ACLs, see <u>Access control list (ACL) overview</u> in the *Amazon Simple Storage Service Developer Guide*.

### Other policy types

Amazon supports additional, less-common policy types. These policy types can set the maximum permissions granted to you by the more common policy types.

- Permissions boundaries A permissions boundary is an advanced feature in which you set the maximum permissions that an identity-based policy can grant to an IAM entity (IAM user or role). You can set a permissions boundary for an entity. The resulting permissions are the intersection of an entity's identity-based policies and its permissions boundaries. Resource-based policies that specify the user or role in the Principal field are not limited by the permissions boundary. An explicit deny in any of these policies overrides the allow. For more information about permissions boundaries, see Permissions boundaries for IAM entities in the IAM User Guide.
- Service control policies (SCPs) SCPs are JSON policies that specify the maximum permissions for an organization or organizational unit (OU) in Amazon Organizations. Amazon Organizations is a service for grouping and centrally managing multiple Amazon Web Services accounts that your business owns. If you enable all features in an organization, then you can apply service control policies (SCPs) to any or all of your accounts. The SCP limits permissions for entities in member accounts, including each Amazon Web Services account root user. For more information about Organizations and SCPs, see How SCPs work in the Amazon Organizations User Guide.
- Session policies Session policies are advanced policies that you pass as a parameter when you programmatically create a temporary session for a role or federated user. The resulting session's permissions are the intersection of the user or role's identity-based policies and the session policies. Permissions can also come from a resource-based policy. An explicit deny in any of these policies overrides the allow. For more information, see Session policies in the IAM User Guide.

### Multiple policy types

When multiple types of policies apply to a request, the resulting permissions are more complicated to understand. To learn how Amazon determines whether to allow a request when multiple policy types are involved, see Policy evaluation logic in the *IAM User Guide*.

### **How Amazon Storage Gateway works with IAM**

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

#### IAM features you can use with Amazon Storage Gateway

| IAM feature                              | Amazon 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 Amazon SGW and other Amazon services work with most IAM features, see Amazon services that work with IAM in the IAM User Guide.

### **Identity-based policies for Amazon 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 Creating IAM policies in the IAM User Guide.

With IAM identity-based policies, you can specify allowed or denied actions and resources as well as the conditions under which actions are allowed or denied. You can't specify the principal in an identity-based policy because it applies to the user or role to which it is attached. To learn about all of the elements that you can use in a JSON policy, see <a href="IAM JSON policy elements reference">IAM JSON policy elements reference</a> in the IAM User Guide.

#### **Identity-based policy examples for Amazon SGW**

To view examples of Amazon SGW identity-based policies, see <u>Identity-based policy examples for</u> Amazon Storage Gateway.

### Resource-based policies within Amazon SGW

| Supports 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 <u>specify a principal</u> in a resource-based policy. Principals can include accounts, users, roles, federated users, or Amazon Web 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 Amazon Web Services 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 How IAM roles differ from resource-based policies in the IAM User Guide.

### **Policy actions for Amazon SGW**

```
Supports policy actions Yes
```

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Action element of a JSON policy describes the actions that you can use to allow or deny access in a policy. Policy actions usually have the same name as the associated Amazon API operation. There are some exceptions, such as *permission-only actions* that don't have a matching API operation. There are also some operations that require multiple actions in a policy. These additional actions are called *dependent actions*.

Include actions in a policy to grant permissions to perform the associated operation.

To see a list of Amazon SGW actions, see <u>Actions Defined by Amazon Storage Gateway</u> in the *Service Authorization Reference*.

Policy actions in Amazon 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 Amazon SGW identity-based policies, see <u>Identity-based policy examples for</u> Amazon Storage Gateway.

### **Policy resources for Amazon SGW**

Supports policy resources Yes

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Resource JSON policy element specifies the object or objects to which the action applies. Statements must include either a Resource or a NotResource element. As a best practice, specify a resource using its <a href="Management-Amazon Resource Name">Amazon Resource Name</a> (ARN). You can do this for actions that support a specific resource type, known as resource-level permissions.

For actions that don't support resource-level permissions, such as listing operations, use a wildcard (\*) to indicate that the statement applies to all resources.

"Resource": "\*"

To see a list of Amazon SGW resource types and their ARNs, see <u>Resources Defined by Amazon Storage Gateway</u> in the *Service Authorization Reference*. To learn with which actions you can specify the ARN of each resource, see <u>Actions Defined by Amazon Storage Gateway</u>.

To view examples of Amazon SGW identity-based policies, see <u>Identity-based policy examples for</u> Amazon Storage Gateway.

### **Policy condition keys for Amazon SGW**

Supports service-specific policy condition keys Yes

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Condition element (or Condition *block*) lets you specify conditions in which a statement is in effect. The Condition element is optional. You can create conditional expressions that use <u>condition operators</u>, such as equals or less than, to match the condition in the policy with values in the request.

If you specify multiple Condition elements in a statement, or multiple keys in a single Condition element, Amazon evaluates them using a logical AND operation. If you specify multiple values for a single condition key, Amazon evaluates the condition using a logical OR operation. All of the conditions must be met before the statement's permissions are granted.

You can also use placeholder variables when you specify conditions. For example, you can grant an IAM user permission to access a resource only if it is tagged with their IAM user name. For more information, see IAM policy elements: variables and tags in the IAM User Guide.

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

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

To view examples of Amazon SGW identity-based policies, see <u>Identity-based policy examples for</u> Amazon Storage Gateway.

#### **ACLs in Amazon SGW**

|--|

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 Amazon SGW

| Partial |
|---------|
|---------|

Attribute-based access control (ABAC) is an authorization strategy that defines permissions based on attributes. In Amazon, these attributes are called *tags*. You can attach tags to IAM entities (users or roles) and to many Amazon resources. Tagging entities and resources is the first step of ABAC. Then you design ABAC policies to allow operations when the principal's tag matches the tag on the resource that they are trying to access.

ABAC is helpful in environments that are growing rapidly and helps with situations where policy management becomes cumbersome.

To control access based on tags, you provide tag information in the <u>condition element</u> 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 <u>What is ABAC?</u> in the *IAM User Guide*. To view a tutorial with steps for setting up ABAC, see <u>Use attribute-based access control</u> (ABAC) in the *IAM User Guide*.

### Using temporary credentials with Amazon SGW

|--|

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

You are using temporary credentials if you sign in to the Amazon Web Services Management Console using any method except a user name and password. For example, when you access Amazon using your company's single sign-on (SSO) link, that process automatically creates temporary credentials. You also automatically create temporary credentials when you sign in to the console as a user and then switch roles. For more information about switching roles, see <a href="Switching to a role">Switching to a role (console)</a> in the IAM User Guide.

You can manually create temporary credentials using the Amazon CLI or Amazon API. You can then use those temporary credentials to access Amazon. Amazon recommends that you dynamically generate temporary credentials instead of using long-term access keys. For more information, see <a href="Temporary security credentials in IAM">Temporary security credentials in IAM</a>.

#### Forward access sessions for Amazon SGW

| Supports forward access sessions (FAS) |
|--|
|--|

When you use an IAM user or role to perform actions in Amazon, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an Amazon Web Service, combined with the requesting Amazon Web Service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other Amazon Web Services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see Forward access sessions.

#### Service roles for Amazon 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 Creating a role to delegate permissions to an Amazon Web Service in the IAM User Guide.



#### Marning

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

#### Service-linked roles for Amazon SGW

Supports service-linked roles

Yes

A service-linked role is a type of service role that is linked to an Amazon Web Service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your Amazon Web Services account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.

For details about creating or managing service-linked roles, see Amazon 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 Amazon Storage Gateway**

By default, users and roles don't have permission to create or modify Amazon SGW resources. They also can't perform tasks by using the Amazon Web Services Management Console, Amazon Command Line Interface (Amazon CLI), or Amazon API. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

To learn how to create an IAM identity-based policy by using these example JSON policy documents, see Creating IAM policies in the IAM User Guide.

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

#### **Topics**

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

### **Policy best practices**

Identity-based policies determine whether someone can create, access, or delete Amazon SGW resources in your account. These actions can incur costs for your Amazon Web Services account. When you create or edit identity-based policies, follow these guidelines and recommendations:

- Get started with Amazon managed policies and move toward least-privilege permissions
  - To get started granting permissions to your users and workloads, use the *Amazon managed policies* that grant permissions for many common use cases. They are available in your Amazon Web Services account. We recommend that you reduce permissions further by defining Amazon customer managed policies that are specific to your use cases. For more information, see <u>Amazon managed policies</u> or <u>Amazon managed policies</u> 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 <a href="Policies and permissions in IAM">Policies and permissions in IAM</a> in the
  IAM User Guide.

- Use conditions in IAM policies to further restrict access You can add a condition to your policies to limit access to actions and resources. For example, you can write a policy condition to specify that all requests must be sent using SSL. You can also use conditions to grant access to service actions if they are used through a specific Amazon Web Service, such as Amazon CloudFormation. For more information, see <a href="IAM JSON policy elements: Condition">IAM JSON policy elements: Condition</a> 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 <a href="IAM Access Analyzer policy validation">IAM Access Analyzer policy validation</a> in the IAM User Guide.
- Require multi-factor authentication (MFA) If you have a scenario that requires IAM users or a
  root user in your Amazon Web Services account, turn on MFA for additional security. To require
  MFA when API operations are called, add MFA conditions to your policies. For more information,
  see Configuring MFA-protected API access in the IAM User Guide.

For more information about best practices in IAM, see <u>Security best practices in IAM</u> in the *IAM User Guide*.

### Using the Amazon SGW console

To access the Amazon Storage Gateway console, you must have a minimum set of permissions. These permissions must allow you to list and view details about the Amazon SGW resources in your Amazon Web Services account. If you create an identity-based policy that is more restrictive than the minimum required permissions, the console won't function as intended for entities (users or roles) with that policy.

You don't need to allow minimum console permissions for users that are making calls only to the Amazon CLI or the Amazon API. Instead, allow access to only the actions that match the API operation that they're trying to perform.

To ensure that users and roles can still use the Amazon SGW console, also attach the Amazon SGW *ConsoleAccess* or *ReadOnly* Amazon 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 Amazon CLI or Amazon API.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "ViewOwnUserInfo",
            "Effect": "Allow",
            "Action": [
                "iam:GetUserPolicy",
                "iam:ListGroupsForUser",
                "iam:ListAttachedUserPolicies",
                "iam:ListUserPolicies",
                "iam:GetUser"
            ],
            "Resource": ["arn:aws-cn:iam::*:user/${aws:username}"]
        },
        {
            "Sid": "NavigateInConsole",
            "Effect": "Allow",
            "Action": [
                "iam:GetGroupPolicy",
                "iam:GetPolicyVersion",
                "iam:GetPolicy",
                "iam:ListAttachedGroupPolicies",
                "iam:ListGroupPolicies",
                "iam:ListPolicyVersions",
                "iam:ListPolicies",
                "iam:ListUsers"
            ],
            "Resource": "*"
        }
    ]
}
```

### **Troubleshooting Amazon Storage Gateway identity and access**

Use the following information to help you diagnose and fix common issues that you might encounter when working with Amazon SGW and IAM.

#### **Topics**

- I am not authorized to perform an action in Amazon SGW
- I am not authorized to perform iam:PassRole
- I want to allow people outside of my Amazon Web Services account to access my Amazon SGW resources

#### I am not authorized to perform an action in Amazon 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-cn: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 Amazon administrator. Your administrator is the person who provided you with your sign-in credentials.

### I am not authorized to perform iam:PassRole

If you receive an error that you're not authorized to perform the iam: PassRole action, your policies must be updated to allow you to pass a role to Amazon SGW.

Some Amazon Web Services allow you to pass an existing role to that service instead of creating a new service role or service-linked role. To do this, you must have permissions to pass the role to the service.

Troubleshooting API Version 2013-06-30 240

The following example error occurs when an IAM user named marymajor tries to use the console to perform an action in Amazon 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-cn:iam::123456789012:user/marymajor is not authorized to perform:
 iam:PassRole

In this case, Mary's policies must be updated to allow her to perform the iam: PassRole action.

If you need help, contact your Amazon administrator. Your administrator is the person who provided you with your sign-in credentials.

### 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 Amazon Identity and Access Management User Guide:

- IAM: Pass an IAM role to a specific Amazon service
- Granting a user permissions to pass a role to an Amazon service
- Available keys for IAM

# I want to allow people outside of my Amazon Web Services account to access my Amazon 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 Amazon SGW supports these features, see <u>How Amazon Storage Gateway</u> works with IAM.

Troubleshooting API Version 2013-06-30 241

- To learn how to provide access to your resources across Amazon Web Services accounts that you own, see <a href="Providing access to an IAM user in another Amazon Web Services account that you own in the IAM User Guide">IAM User Guide</a>.
- To learn how to provide access to your resources to third-party Amazon Web Services accounts, see <u>Providing access to Amazon Web Services accounts owned by third parties</u> in the *IAM User Guide*.
- To learn how to provide access through identity federation, see <u>Providing access to externally</u> authenticated users (identity federation) in the *IAM User Guide*.
- To learn the difference between using roles and resource-based policies for cross-account access, see How IAM roles differ from resource-based policies in the IAM User Guide.

### Using tags to control access to your gateway and resources

To control access to gateway resources and actions, you can use Amazon 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 <a href="Controlling Access Using Tags">Controlling Access Using Tags</a>.

#### **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.

```
"storagegateway:ListTagsForResource",
                         "storagegateway:ListFileShares",
                         "storagegateway:DescribeNFSFileShares"
                     ],
                     "Resource": "*",
                     "Condition": {
                         "StringEquals": {
                             "aws:ResourceTag/allowListAndDescribe": "yes"
                         }
                     }
      },
          "Effect": "Allow",
          "Action": [
               "storagegateway: *"
          ],
          "Resource": "arn:aws:storagegateway:region:account-id:*/*"
      }
  ]
}
```

### **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.

```
{
    "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.
- 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 <a href="CreateSMBFileShare"><u>CreateSMBFileShare</u></a> API operation. Or you can turn on Windows ACLs on an existing SMB file share by using the <a href="UpdateSMBFileShare">UpdateSMBFileShare</a> 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

- Create a file gateway if you don't already have one. For more information, see <u>Creating Your Gateway</u>.
- 2. If the gateway is not joined to a domain, add it to a domain. For more information, see <u>Using</u>
  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 <u>AdminUsersList</u>, 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.

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.

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

- Choose the file share and choose **Edit file share**.
- For the File/directory access controlled by option, choose Windows Access Control List. 2.
- 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.
- 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.

### Logging and monitoring in Amazon Storage Gateway

Storage Gateway is integrated with Amazon CloudTrail, a service that provides a record of actions taken by a user, role, or an Amazon 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 Amazon CloudTrail User Guide.

### Storage Gateway information in CloudTrail

CloudTrail is activated on your Amazon account when you create the account. When activity occurs in Storage Gateway, that activity is recorded in a CloudTrail event along with other Amazon service events in **Event history**. You can view, search, and download recent events in your Amazon account. For more information, see Viewing Events with CloudTrail Event History.

For an ongoing record of events in your Amazon account, including events for 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 Amazon Regions. The trail logs events from

Logging and monitoring API Version 2013-06-30 248 all Regions in the Amazon partition and delivers the log files to the Amazon S3 bucket that you specify. Additionally, you can configure other Amazon services to further analyze and act upon the event data collected in CloudTrail logs. For more information, see the following:

- 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 <u>Actions</u> 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 Amazon Identity and Access Management (IAM) user credentials.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another Amazon service.

For more information, see the CloudTrail userIdentity Element.

### **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.

```
"type": "IAMUser",
                "principalId": "AIDAII5AUEPBH2M7JTNVC",
                "arn": "arn:aws:iam::111122223333:user/StorageGateway-team/JohnDoe",
                "accountId": "111122223333",
                "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
                 "userName": "JohnDoe"
               },
                  "eventTime": "2014-12-04T16:19:00Z",
                  "eventSource": "storagegateway.amazonaws.com",
                  "eventName": "ActivateGateway",
                  "awsRegion": "us-east-2",
                  "sourceIPAddress": "192.0.2.0",
                  "userAgent": "aws-cli/1.6.2 Python/2.7.6 Linux/2.6.18-164.el5",
                   "requestParameters": {
                                            "gatewayTimezone": "GMT-5:00",
                                            "gatewayName": "cloudtrailgatewayvtl",
                                            "gatewayRegion": "us-east-2",
                                            "activationKey": "EHFBX-1NDD0-P0IVU-PI259-
DHK88",
                                            "gatewayType": "VTL"
                                                 },
                                                 "responseElements": {
                                                                        "gatewayARN":
 "arn:aws:storagegateway:us-east-2:111122223333:gateway/cloudtrailgatewayvtl"
                                                 "requestID":
 "54BTFGNQI71987UJD2IHTCT8NF1Q8GLLE1QEU3KPGG6F0KSTAUU0",
                                                 "eventID": "635f2ea2-7e42-45f0-
bed1-8b17d7b74265",
                                                 "eventType": "AwsApiCall",
                                                 "apiVersion": "20130630",
                                                 "recipientAccountId": "444455556666"
             }]
}
```

The following example shows a CloudTrail log entry that demonstrates the ListGateways action.

```
"arn": "arn:aws:iam::111122223333:user/StorageGateway-
team/JohnDoe",
                                 "accountId:" 111122223333", " accessKeyId ":"
 AKIAIOSFODNN7EXAMPLE",
                                 " userName ":" JohnDoe "
                                },
                                 " eventTime ":" 2014 - 12 - 03T19: 41: 53Z ",
                                 " eventSource ":" storagegateway.amazonaws.com ",
                                 " eventName ":" ListGateways ",
                                 " awsRegion ":" us-east-2 ",
                                 " sourceIPAddress ":" 192.0.2.0 ",
                                 " userAgent ":" aws - cli / 1.6.2 Python / 2.7.6
 Linux / 2.6.18 - 164.el5 ",
                                 " requestParameters ":null,
                                 " responseElements ":null,
                                 "requestID ":"
 6U2N42CU37KA08BG6V1I23FRSJ1Q8GLLE1QEU3KPGG6F0KSTAUU0 ",
                                 " eventID ":" f76e5919 - 9362 - 48ff - a7c4 -
 d203a189ec8d ",
                                 " eventType ":" AwsApiCall ",
                                 " apiVersion ":" 20130630 ",
                                 " recipientAccountId ":" 444455556666"
              }]
}
```

### **Compliance validation for Amazon Storage Gateway**

Third-party auditors assess the security and compliance of Amazon Storage Gateway as part of multiple Amazon compliance programs. These include SOC, PCI, ISO, FedRAMP, HIPAA, MTCS, C5, K-ISMS, ENS High, OSPAR, and HITRUST CSF.

For a list of Amazon services in scope of specific compliance programs, see <u>Amazon Services in Scope by Compliance Program</u>. For general information, see <u>Amazon Compliance Programs</u>.

You can download third-party audit reports using Amazon Artifact. For more information, see Downloading Reports in Amazon 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. Amazon provides the following resources to help with compliance:

Compliance validation API Version 2013-06-30 251

- <u>Security and Compliance Quick Start GuidesSecurity and Compliance Quick Start Guides</u> These deployment guides discuss architectural considerations and provide steps for deploying security-and compliance-focused baseline environments on Amazon.
- <u>Architecting for HIPAA Security and Compliance Whitepaper</u> This whitepaper describes how companies can use Amazon to create HIPAA-compliant applications.
- <u>Amazon Compliance Resources</u> This collection of workbooks and guides might apply to your industry and location.
- <u>Evaluating resources with rules</u> in the *Amazon Config Developer Guide* The Amazon Config service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- <u>Amazon Security Hub</u> This Amazon service provides a comprehensive view of your security state within Amazon that helps you check your compliance with security industry standards and best practices.

### Resilience in Amazon Storage Gateway

The Amazon global infrastructure is built around Amazon Regions and Availability Zones. Amazon Regions provide multiple physically separated and isolated Availability Zones, which are connected with low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For more information about Amazon Regions and Availability Zones, see <u>Amazon Global</u> Infrastructure.

In addition to the Amazon 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 <u>Using VMware vSphere High Availability with Storage Gateway</u>.

### Infrastructure security in Amazon Storage Gateway

As a managed service, Amazon Storage Gateway is protected by the Amazon global network security procedures that are described in Security Pillar - Amazon Well-Architected Framework.

Resilience API Version 2013-06-30 252

You use Amazon 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 Amazon Security Token Service (Amazon STS) to generate temporary security credentials to sign requests.



#### Note

You should treat the Amazon 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.

Amazon reviews, analyzes, and remediates CVEs on a regular basis, which we incorporate into the Amazon Storage Gateway in our normal software release cycle. These updates are typically released twice per year and are applied as part of the normal gateway update process during scheduled maintenance windows.

### **Amazon Security Best Practices**

Amazon 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 Amazon Security Best Practices.

### **Troubleshooting**

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: on-premises gateway issues
- Troubleshooting: Microsoft Hyper-V setup
- Troubleshooting: Amazon EC2 gateway issues
- Troubleshooting: hardware appliance issues
- Troubleshooting: File Gateway issues
- Troubleshooting: file share issues
- High Availability Health Notifications
- Troubleshooting: high availability issues

### 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 Amazon Web Services Support to connect to your gateway to assist with troubleshooting.

The following table lists typical issues that you might encounter working with your on-premises gateways.

| Issue   | Action to Take   |
|---|--|
| You cannot find the IP address of your gateway. | Use the hypervisor client to connect to your host to find the gateway IP address.  |
|   | <ul> <li>For VMware ESXi, the VM's IP address can be found in the<br/>vSphere client on the <b>Summary</b> tab.</li> </ul> |

| Issue                                       | Action to Take   |
|---|--|
|   | <ul> <li>For Microsoft Hyper-V, the VM's IP address can be found by<br/>logging into the local console.</li> </ul>   |
|   | If you are still having trouble finding the gateway IP address:  |
|   | <ul> <li>Check that the VM is turned on. Only when the VM is turned on<br/>does an IP address get assigned to your gateway.</li> </ul>   |
|   | <ul> <li>Wait for the VM to finish startup. If you just turned on your VM,<br/>then it might take several minutes for the gateway to finish its<br/>boot sequence.</li> </ul>  |
| You're having network or firewall problems. | <ul> <li>Allow the appropriate ports for your gateway.</li> <li>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 Amazon. For more information about network and firewall requirements, see <a href="Network and firewall requirements">Network and firewall requirements</a>.</li> </ul> |

## Amazon Storage Gateway **Action to Take** Issue Your gateway's activatio n fails when you click the from your client. **Proceed to Activation** button in the Storage **Gateway Management** Console. connectivity. You need to improve bandwidth between your

- Check that the gateway VM can be accessed by pinging the VM
- 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
- 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.

gateway and Amazon.

You can improve the bandwidth from your gateway to Amazon by setting up your internet connection to Amazon 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 Amazon and you want to avoid bandwidth contention, especially during a snapshot restore. For high-throughput workload needs, you can use Amazon Direct Connect to establish a dedicated network connection between your on-premises gateway and Amazon. To measure the bandwidth of the connection from your gateway to Amazon, use the CloudBytesDownloaded and CloudBytesUploaded metrics of the gateway. For more on this subject, see Performan ce and optimization. Improving your internet connectivity helps to ensure that your upload buffer does not fill up.

### **Action to Take** Issue Throughput to or from On the Gateway tab of the Storage Gateway console, verify that the IP addresses for your gateway VM are the same that you your gateway drops to see using your hypervisor client software (that is, the VMware zero. 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 Amazon as described in Testing your gateway's network connectivity. Check your gateway's network adapter configuration in your hypervisor management client and ensure that all the interfaces you intend to use for the gateway are activated. • Check your gateway's network adapter configuration in the gateway local console. For instructions, see Configuring your gateway network settings. You can view the throughput to and from your gateway from the Amazon CloudWatch console. For more information about measuring throughput to and from your gateway to Amazon, see Performance and optimization. See Troubleshooting: Microsoft Hyper-V setup, which discusses You are having trouble importing (deploying) some of the common issues of deploying a gateway on Microsoft Storage Gateway on Hyper-V. 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 Amazon". | 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 Amazon Web Services Support. |

# Turning on Amazon Web Services 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 Amazon Web Services Support to access your gateway to assist you with troubleshooting gateway issues. By default, Amazon Web Services Support access to your gateway is turned off. You turn on this access through the host's local console. To give Amazon Web Services 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 Amazon Web Services Support access to your gateway

- 1. Log in to your host's local console.
  - VMware ESXi for more information, see <u>Accessing the Gateway Local Console with VMware ESXi</u>.
  - Microsoft Hyper-V for more information, see <u>Access the Gateway Local Console with</u> Microsoft Hyper-V.
- 2. At the prompt, enter the corresponding numeral to select **Gateway Console**.
- 3. Enter **h** to open the list of available commands.
- 4. Do one of the following:
  - If your gateway is using a public endpoint, in the AVAILABLE COMMANDS window, enter
     open-support-channel to connect to customer support for Storage Gateway. Allow TCP
     port 22 so you can open a support channel to Amazon. 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 Amazon. 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.

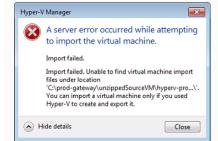
- After the support channel is established, provide your support service number to Amazon Web Services Support so Amazon Web Services Support can provide troubleshooting assistance.
- When the support session is completed, enter  $\mathbf{q}$  to end it. Don't close the session until Amazon Web Services Support notifies you that the support session is complete.
- Enter **exit** to log out of the Storage Gateway console.
- Follow the prompts to exit the local console.

### Troubleshooting: Microsoft Hyper-V setup

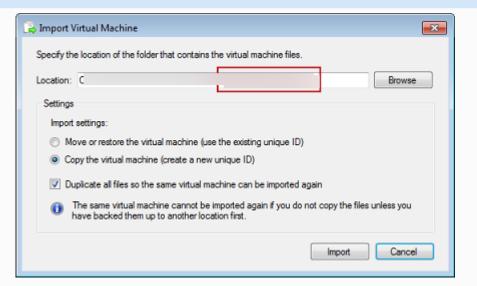
The following table lists typical issues that you might encounter when deploying Storage Gateway on the Microsoft Hyper-V platform.

| Issue  | Action to Take   |
|--|--|
| You try to import a gateway and receive the error message: "Import failed. Unable to find virtual machine import file under location". | <ul> <li>This error can occur for the following reasons:</li> <li>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, as the following example shows:</li> </ul> |

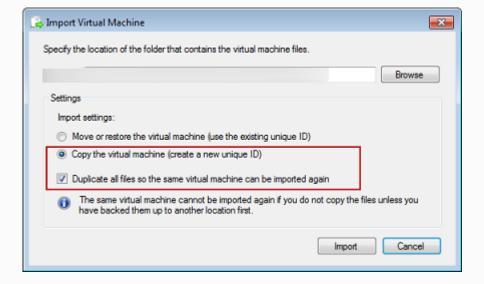
#### Issue



#### **Action to Take**

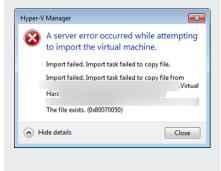


• If you have already deployed a gateway and you did not select the Copy the virtual machine option and check the Duplicate all files option in the Import Virtual Machine dialog box, then the VM was created in the location where you have the unzipped gateway files and you cannot import from this location 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. The following example shows the options that you must check if you plan on creating multiple gateways from one unzipped source files location.



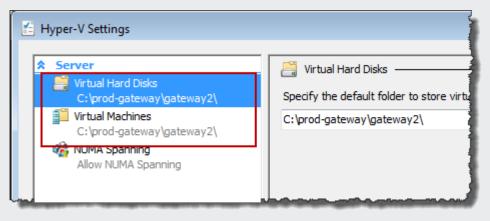
#### Issue

You try to import a gateway and receive the error message: "Import failed. Import task failed to copy file."

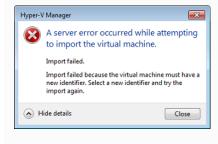


#### **Action to Take**

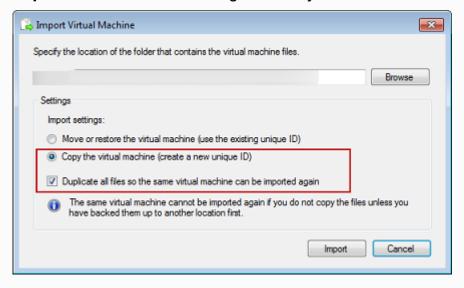
If you have already deployed a gateway and you try to reuse the default folders that store the virtual hard disk files and virtual machine configuration files, then this error will occur. To fix this problem, specify new locations in the **Hyper-V Settings** dialog box.



You try to import a gateway and receive an error message: "Import failed. Import failed because the virtual machine must have a new identifier. Select a new identifier and try the import again."

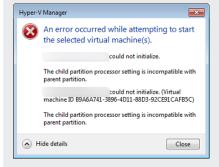


When you import the gateway make sure you select the **Copy the virtual machine** option and check the **Duplicate all files** option in the **Import Virtual Machine** dialog box to create a new unique ID for the VM. The following example shows the options in the **Import Virtual Machine** dialog box that you should use.



#### Issue

You try to start a gateway VM and receive an error message "The child partition processor setting is incompatible with parent partition."

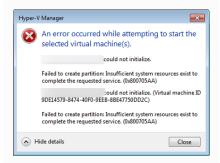


#### **Action to Take**

This error is likely caused by a CPU discrepancy between the required CPUs for the gateway and the available CPUs on the host. Ensure that the VM CPU count is supported by the underlying hypervisor.

For more information about the requirements for Storage Gateway, see File Gateway setup requirements.

You try to start a gateway VM and receive an error message "Failed to create partition: Insuffici ent resources exist to complete the requested service."



This error is likely caused by a RAM discrepancy between the required RAM for the gateway and the available RAM on the host.

For more information about the requirements for Storage Gateway, see File Gateway setup requirements.

Your snapshots and gateway software updates are occurring at slightly different times than expected.

The gateway VM's clock might be offset from the actual time, known as clock drift. Check and correct the VM's time using local gateway console's time synchronization option. For more information, see <a href="Configuring a Network Time Protocol">Configuring a Network Time Protocol</a> (NTP) server for your gateway.

#### **Action to Take** Issue You need to put the Access the host as you do a typical Microsoft Windows server. For unzipped Microsoft Hyperexample, if the hypervisor host is name hyperv-server, then V Storage Gateway files on you can use the following UNC path \\hyperv-server\c\$ the host file system. which assumes that the name hyperv-server can be resolved or is defined in your local hosts file. You are prompted Add your user credentials as a local administrator for the for credentials when hypervisor host by using the Sconfig.cmd tool. connecting to hypervisor. Your credentials did not work Your system administrator does not allow the use of default credentials to log on to the remote computer HYPERV-SERVER because its identity is not fully verified. Please enter new credentials. Use another account Remember my credentials The logon attempt failed OK Cancel You may notice poor For information about a workaround, see the Microsoft network performance documentation, see Poor network performance on virtual if you turn on virtual machines on a Windows Server 2012 Hyper-V host if VMQ is machine queue (VMQ) turned on. for a Hyper-V host that's using a Broadcom network

### 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 onpremises gateway and a gateway deployed in Amazon EC2, see <a href="Deploying an Amazon EC2">Deploying an Amazon EC2</a> instance to host an S3 File Gateway.

For information about using ephemeral storage, see <u>Using ephemeral storage with EC2 gateways</u>.

#### **Topics**

adapter.

- · 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 Amazon Web Services 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 <a href="Adding a security group rule">Adding a security group rule</a> in the Amazon EC2 User Guide for Linux Instances.
- 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 Amazon Web Services Support to help troubleshoot your **Amazon EC2 gateway**

Storage Gateway provides a local console you can use to perform several maintenance tasks, including allowing Amazon Web Services Support to access your gateway to assist you with troubleshooting gateway issues. By default, Amazon Web Services 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 Amazon Web Services 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 Amazon Web Services Support access for a gateway deployed on an Amazon EC2 instance

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.

- At the prompt, enter 6 Command Prompt to open the Amazon Web Services Support Channel console.
- Enter **h** to open the **AVAILABLE COMMANDS** window. 3.
- 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 Amazon. 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 Amazon. 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.

After the support channel is established, provide your support service number to Amazon Web Services Support so Amazon Web Services 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

The following topics discuss issues that you might encounter with the Amazon 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 Amazon 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 <u>iDRAC9 Virtual Power Cycle: Remotely power cycle Dell EMC PowerEdge Servers</u> on the Dell Technologies InfoHub website.

### Where do you obtain Dell iDRAC support?

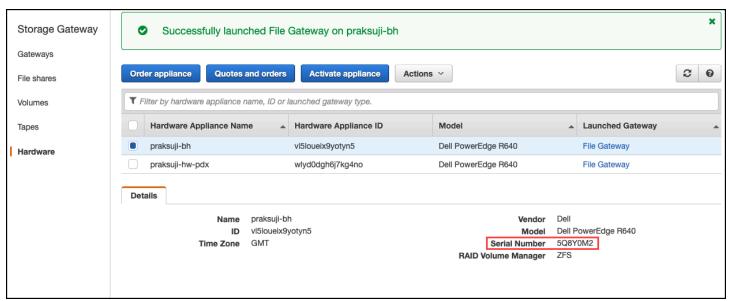
The Dell PowerEdge R640 server comes with the Dell iDRAC management interface. We recommend the following:

- If you use the iDRAC management interface, you should change the default password. For more
  information about the iDRAC credentials, see <u>Dell PowerEdge What is the default sign-in</u>
  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

To find the serial number of the hardware appliance, go to the **Hardware appliance overview** page in the Storage Gateway console, as shown following.

Storage Gateway console hardware tab with appliance selected and details shown.



Storage Gateway console hardware tab with appliance selected and details shown.

### Where to obtain hardware appliance support

To contact Amazon about technical support for your hardware appliance, see <u>Amazon Web Services</u> Support.

The Amazon Web Services 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 Amazon

- 1. Open the hardware console.
- Choose Open Support Channel as shown following.
   hardware appliance console with support channel status shown.



hardware appliance console with support channel status shown.

The assigned port number should appear within 30 seconds, if there are no network connectivity or firewall issues.

3. Note the port number and provide it to Amazon Web Services 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: Using CloudWatch metrics

### Error: 1344 (0x00000540)

While migrating files to Amazon S3 you may encounter an ERROR 1344 ( $0 \times 00000540$ ) 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 Amazon Storage Gateway servers. Clock synchronization issues can negatively impact connectivity between the gateway and Amazon. 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: 1344 (0x00000540) API Version 2013-06-30 270

### **Error: InaccessibleStorageClass**

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

Here, usually your File Gateway encounters the error when it tries to either upload the specified object to S3 bucket or read the object from S3 bucket. With this error, generally the object has moved to Amazon S3 Glacier and is in either the S3 Glacier or S3 Glacier Deep Archive storage class.

#### To resolve an InaccessibleStorageClass error

Move the object from the S3 Glacier or S3 Glacier Deep Archive storage class back to S3.

If you move the object to the S3 bucket to fix an upload error, the file is eventually uploaded. If you move the object to the S3 bucket 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 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.

#### To resolve an InvalidObjectState error

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

- 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 Amazon Web Services Management Console or Amazon CLI.
- 2. Delete the file in Amazon S3 using the Amazon Web Services Management Console or Amazon 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.

#### 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 IAMrole. For information about IAM role, see <u>Tutorial</u>: <u>delegate</u> <u>access across Amazon accounts using IAM roles</u>.

### **Error: S3AccessDenied**

You can get an S3AccessDenied error for a file share's Amazon S3 bucket access Amazon 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.

#### To resolve an S3AccessDenied error

Error: ObjectMissing API Version 2013-06-30 272

• 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 low. When available root disk space is less than 1 GB, the gateway stops generating detailed CloudWatch log notifications as a precautionary measure.

#### To resolve a DroppedNotifications error

- 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.
- Increase the size of the gateway's root storage disk if available space is less than 1 GB. Refer to 2. 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 Amazon Storage Gateway Hardware Appliance.

3. Restart your gateway.

Error: DroppedNotifications API Version 2013-06-30 273

### **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 <u>maintenance start time</u>, 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: 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:

Notification: HardReboot API Version 2013-06-30 274

- If the IndexFetch metric is greater than 0 when you run an 1s 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 Amazon Web Services Support.

Discuss with Amazon Web Services 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.amazonaws.cn/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 <a href="Create a bucket">Create a bucket</a> 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.
- 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 <a href="Granting">Granting</a> 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 <a href="Your file">Your file</a> share is stuck in CREATING status, preceding.
- 2. If the S3 bucket exists, then verify that Amazon 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 Amazon Security Token Service, see <a href="Activating">Activating</a> and deactivating Amazon STS in an Amazon 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.

Can't create a file share API Version 2013-06-30 278

### 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 <a href="Best Practices for File Gateway">Best Practices for File Gateway</a>.

### 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 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 <u>Using Identity-Based Policies</u> (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 Amazon 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 Amazon 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 Amazon 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 objects in your Amazon S3 bucket.

#### 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 Amazon Command Line Interface (Amazon CLI):

- To refresh the cache using the Storage Gateway console, see Refreshing objects in your Amazon S3 bucket.
- To refresh the cache using the Amazon CLI:
  - 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 <a href="How can I automate the RefreshCache operation">How can I automate the RefreshCache operation</a> 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 <u>Troubleshooting</u>: 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 <u>maintenance start</u> <u>time</u>, 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

Health notifications API Version 2013-06-30 284

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 Amazon Web Services 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.

Metrics API Version 2013-06-30 285

# **Best Practices for File Gateway**

Follow the best practices in this section when you create gateways, file shares, and snapshots.

### **Topics**

- Best practices: recovering your data
- Best practices: managing multipart uploads
- 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
- Allowing specific NFS clients to mount your file share

### 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.

#### **Topics**

- Recovering from an unexpected virtual machine shutdown
- Recovering your data from a malfunctioning cache disk
- Recovering your data from an inaccessible data center

### 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 <u>Testing your gateway</u>'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.

- Create and activate a new File Gateway on an Amazon EC2 host. For more information, see
   Deploying an Amazon EC2 instance to host an 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 <u>Mount and use your file share</u>.

### 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 <a href="Uploading and copying objects using multipart">Upload in the Amazon Simple Storage Service User Guide</a>.

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 <a href="Configuring">Configuring</a> a bucket lifecycle configuration to delete incomplete multipart uploads in the Amazon Simple Storage Service User Guide.

### 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 <u>robocopy</u> and <u>xcopy</u> 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 Amazon.

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.

### 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:::TestBucket/\*" bucket.

```
{
   "Version": "2012-10-17",
   "Statement":[
      {
          "Sid": "DenyMultiWrite",
          "Effect": "Deny",
          "Principal": "*",
          "Action":[
             "s3:DeleteObject",
             "s3:PutObject"
          "Resource":"arn:aws:s3:::TestBucket/*",
          "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.

# Allowing specific NFS clients to mount your file share

We recommend that you change the allowed Network File System (NFS) client settings for your file share. If you don't, any client on your network can mount your file share. For information about how to edit your NFS client settings, see Editing access settings for your NFS file share.

### **Additional Storage Gateway resources**

This section describes Amazon and third-party software, tools, and resources that can help you set up or manage your gateway, and also Storage Gateway quotas.

#### **Topics**

- Host setup
- Getting an activation key for your gateway
- Support for file attributes in Amazon S3 File Gateway
- Using Amazon Direct Connect with Storage Gateway
- Port Requirements
- Active Directory service account permission requirements
- Connecting to Your Gateway
- Understanding Storage Gateway Resources and Resource IDs
- Tagging Storage Gateway resources
- Working with open-source components for Amazon Storage Gateway
- Quotas
- Using storage classes
- Using Kubernetes Container Storage Interface drivers
- Amazon Storage Gateway Terraform module

### **Host setup**

The following topics provide information about setting up the virtual machine host platform for your gateway.

### **Topics**

- Configuring VMware for Storage Gateway
- Synchronizing Your Gateway VM Time
- Deploying an Amazon EC2 instance to host an S3 File Gateway
- Deploy an Amazon EC2 with Default Settings
- Modify Amazon EC2 instance metadata options

Host setup API Version 2013-06-30 292

### **Configuring VMware for Storage Gateway**

When configuring VMware for Storage Gateway, make sure to synchronize your VM time with your host time, configure VM to use paravirtualized disk controllers when provisioning storage and provide protection from failures in the infrastructure layer supporting a gateway VM.

### **Topics**

- Synchronizing VM Time with Host Time
- Using Storage Gateway with VMware High Availability

### **Synchronizing VM Time with 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.



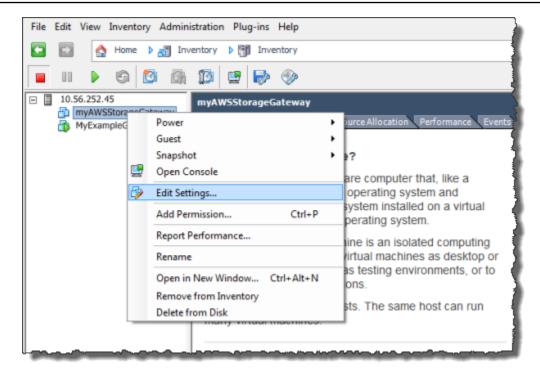
#### 

Synchronizing the VM time with the host time is required for successful gateway activation.

### To synchronize VM time with host time

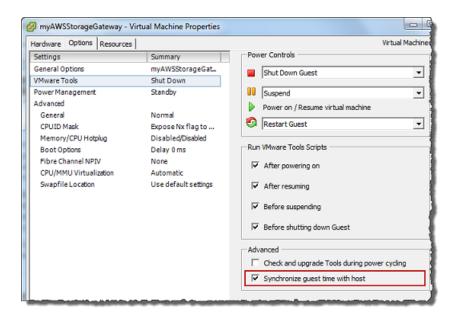
- Configure your VM time.
  - In the vSphere client, open the context (right-click) menu for your gateway VM, and choose **Edit Settings**.

The Virtual Machine Properties dialog box opens.



- b. Choose the **Options** tab, and choose **VMware Tools** in the options list.
- c. Check the **Synchronize guest time with host** option, and then choose **OK**.

The VM synchronizes its time with the host.

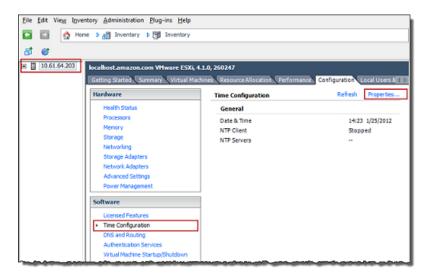


2. Configure the host time.

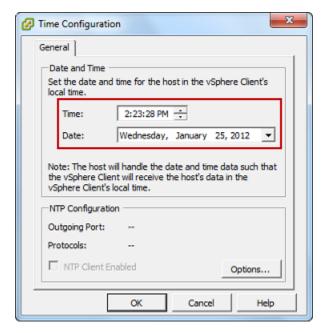
It is important to make sure that your host clock is set to the correct time. If you have not configured your host clock, perform the following steps to set and synchronize it with an NTP server.

- a. In the VMware vSphere client, select the vSphere host node in the left pane, and then choose the **Configuration** tab.
- b. Select **Time Configuration** in the **Software** panel, and then choose the **Properties** link.

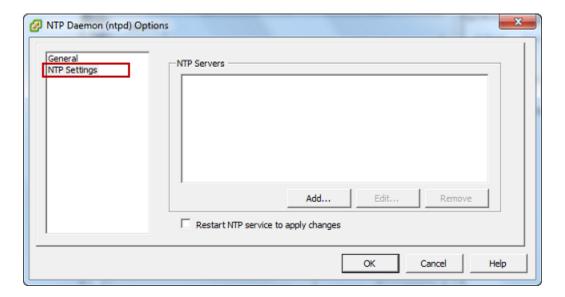
The **Time Configuration** dialog box appears.



c. In the **Date and Time** panel, set the date and time.



- 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 pane.



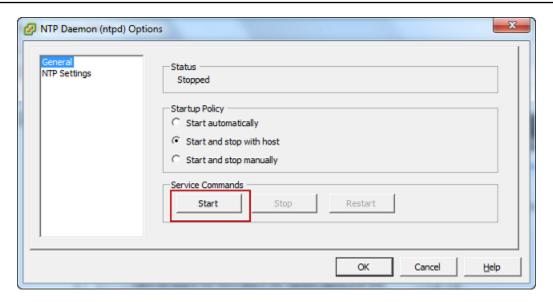
- ii. Choose Add to add a new NTP server.
- iii. In the **Add NTP Server** dialog box, type the IP address or the fully qualified domain name of an NTP server, and then choose **OK**.

You can use pool.ntp.org as shown in the following example.



- iv. In the NTP Daemon (ntpd) Options dialog box, choose General in the left pane.
- v. In the **Service Commands** pane, 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.

### **Using Storage Gateway with VMware High Availability**

VMware High Availability (HA) is a component of vSphere that can provide protection from failures in the infrastructure layer supporting a gateway VM. VMware HA does this by using multiple hosts configured as a cluster so that if a host running a gateway VM fails, the gateway VM can be restarted automatically on another host within the cluster. For more information about VMware HA, see VMware HA: Concepts and Best Practices on the VMware website.

To use Storage Gateway with VMware HA, we recommend doing the following things:

- 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.

### **Synchronizing Your Gateway VM Time**

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 <a href="Synchronizing VM">Synchronizing VM</a> <a href="Time with Host Time">Time with Host Time</a>. For a gateway deployed on Microsoft Hyper-V or Linux KVM, you should periodically check your VM's time using the procedure described following.

To view and synchronize the time of a hypervisor gateway VM 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 <u>Access the</u> Gateway Local Console with Microsoft Hyper-V.
  - For more information on logging in to the local console for Linux Kernel-based Virtuam Machine (KVM), see Accessing the Gateway Local Console with Linux KVM.
- 2. On the **Storage Gateway Configuration** main menu, enter **4** for **System Time Management**.

```
Storage Gateway Configuration
Currently connected network adapters:
##
## eth0: 10.0.0.45
1: SOCKS Proxy Configuration
2: Network Configuration
3: Test Network Connectivity
4: System Time Management
5: Gateway Console
6: View System Resource Check (0 Errors)
0: Stop
        Storage Gateway
Press "x" to exit session
Enter command: _
```

3. On the System Time Management menu, enter 1 for View and Synchronize System Time.

```
System Time Management

1: View and Synchronize System Time

Press "x" to exit

Enter command: _
```

If the result indicates that you should synchronize your VM's time to the NTP time, enter y.
 Otherwise, enter n.

If you enter **y** to synchronize, the synchronization might take a few moments.

The following screen shot shows a VM that doesn't require time synchronization.

```
System Time Management

1: View and Synchronize System Time

Press "x" to exit

Enter command: 1

Current System Time: Sat Aug 22 00:33:41 UTC 2015

Determining current NTP time (this may take a few seconds ...)

Your Storage Gateway UM system time differs from NTP time by 0.217617 seconds

A sync is recommended if the time differs by more than 60 seconds

Do you want to sync Storage Gateway UM system time with NTP time? [y/n]: _
```

The following screen shot shows a VM that does require time synchronization.

```
System Time Management

1: View and Synchronize System Time

Press "x" to exit

Enter command: 1

Current System Time: Sat Aug 22 00:33:41 UTC 2015

Determining current NTP time (this may take a few seconds ...)

Your Storage Gateway UM system time differs from NTP time by 61.217617 seconds

A sync is recommended if the time differs by more than 60 seconds

Do you want to sync Storage Gateway UM system time with NTP time? [y/n]: __
```

### Deploying an Amazon EC2 instance to host an S3 File Gateway

You can deploy and activate an Amazon S3 File Gateway on an Amazon Elastic Compute Cloud (Amazon EC2) instance. The Amazon Storage Gateway Amazon Machine Image (AMI) is available as a community AMI.

### Note

Storage Gateway community AMIs are published and fully supported by Amazon. You can see that the publisher is Amazon, a verified provider.

#### To deploy an Amazon EC2 instance to host your Amazon S3 File Gateway

- Start setting up a new gateway using the Storage Gateway console. For instructions, see <u>Set up an Amazon S3 File Gateway</u>. 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 Amazon 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.
- In the **Instance type** section, for **Instance type**, choose the hardware configuration for your 4. 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 for Linux Instances.



#### 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.

- In the **Key pair (login)** section, for **Key pair name required**, select the key pair you want to 5. 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.
- In the **Network settings** section, review the preconfigured settings and choose **Edit** to make changes to the following fields:
  - For **VPC required**, choose the VPC where you want to launch your Amazon EC2 instance. a. For more information, see How Amazon VPC works in the Amazon Virtual Private Cloud User Guide.
  - (Optional) For **Subnet**, choose the subnet where you want to launch your Amazon EC2 instance.
  - For **Auto-assign Public IP**, choose **Enable**. c.

- In the Firewall (security groups) subsection, review the preconfigured settings. You can 7. 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.
- 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 onetime 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.

- 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 Amazon 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 <u>Set up an Amazon S3 File Gateway</u>. When you reach the **Platform options** section, choose **Amazon EC2** as the **Host platform**, then choose **Launch instance** to open the Amazon 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 Amazon Region in the URL.

Query the Systems Manager parameter store. You can use the Amazon 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 Amazon Web Services 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"
}
```

}

### **Deploy an Amazon EC2 with Default Settings**

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 Amazon Storage Gateway Amazon Machine Image (AMI) is available as a community AMI.



#### Note

Storage Gateway community AMIs are published and fully supported by Amazon. You can see that the publisher is Amazon, a verified provider.

- 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.
- Select Launch instance to open the Amazon Storage Gateway AMI template in the Amazon EC2 console and customize additional settings such as Instance types, Network settings and Configure storage.
- 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<br>System<br>Protocol       |
|------|----------|----------------------------------|
| 80   | ТСР      | HTTP<br>access for<br>activation |
| 111  | TCP, UDP | NFSv3                            |
| 139  | TCP, UDP | SMB                              |
| 445  | ТСР      | SMB                              |
| 2049 | TCP, UDP | NFS                              |

| Port  | Protocol | File<br>System<br>Protocol |  |  |
|-------|----------|----------------------------|--|--|
| 20048 | TCP, UDP | NFSv3                      |  |  |

#### Configure storage

| Default<br>Settings          | AMI Root<br>Volume | Volume 2<br>Cache |
|------------------------------|--------------------|-------------------|
| Device<br>Name               |                    | '/dev/sdb'        |
| Size                         | 80 Gib             | 165 GiB           |
| Volume<br>Type               | gp3                | gp3               |
| IOPS                         | 3000               | 3000              |
| Delete on<br>terminati<br>on | Yes                | Yes               |
| Encrypted                    | No                 | No                |
| Throughpu<br>t               | 125                | 125               |

### **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 <a href="How Instance">How Instance</a> 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.

### 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 Amazon 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 guery 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 Amazon 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**

**Getting Activation Key** API Version 2013-06-30 307

- 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 Amazon General Reference Guide.
- vpc\_endpoint The VPC endpoint name for your gateway, for example
  vpce-050f90485f28f2fd0-iep0e8vq.storagegateway.uswest-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 (curl) API Version 2013-06-30 308

### 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) {
```

Linux (bash/zsh) API Version 2013-06-30 309

### 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. If you are connecting to your Amazon EC2 instance from a Windows computer, log in as *admin*.
- After you log in and see the Amazon 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 Amazon 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 <a href="File Attribute Constants">File Attribute Constants</a> 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.

Using your local console API Version 2013-06-30 310



#### 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 Amazon Direct Connect with Storage Gateway**

Amazon Direct Connect links your internal network to the Amazon Web Services Cloud. By using Amazon 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 Amazon.

Storage Gateway uses public endpoints. With an Amazon 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 Amazon Region as the Amazon Direct Connect location, or it can be in a different Amazon Region.

The following illustration shows an example of how Amazon Direct Connect works with Storage Gateway.

network architecture showing Storage Gateway connected to the cloud using Amazon direct connect.

The following procedure assumes that you have created a functioning gateway.

### To use Amazon Direct Connect with Storage Gateway

- Create and establish an Amazon Direct Connect connection between your on-premises data center and your Storage Gateway endpoint. For more information about how to create a connection, see <u>Getting Started with Amazon Direct Connect</u> in the *Amazon Direct Connect User Guide*.
- 2. Connect your on-premises Storage Gateway appliance to the Amazon Direct Connect router.
- 3. Create a public virtual interface, and configure your on-premises router accordingly. For more information, see <a href="Creating a Virtual Interface">Creating a Virtual Interface</a> in the Amazon Direct Connect User Guide.

For details about Amazon Direct Connect, see <u>What is Amazon Direct Connect?</u> in the *Amazon Direct Connect User Guide*.

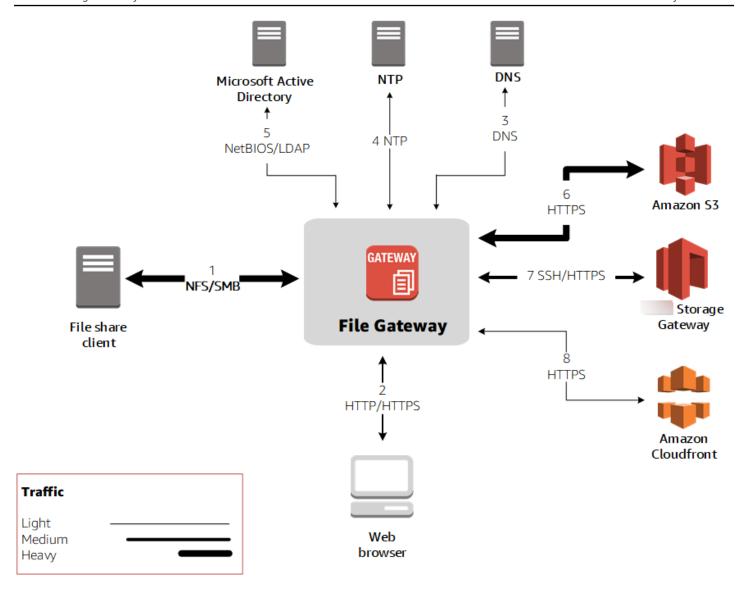
### **Port Requirements**

Storage Gateway requires the following ports for its operation. Some ports are common to all gateway types and are required by all gateway types. Other ports are required by specific gateway types. In this section, you can find an illustration of the required ports and a list of the ports required by each gateway type.

#### **File Gateways**

The following illustration shows the ports to open for File Gateways' operation.

Port Requirements API Version 2013-06-30 312



The following ports are common to all gateway types and are required by all gateway types.

| From                  | То                     | Protocol                                      | Port        | How Used  |
|-----------------------|------------------------|---|-------------|---|
| Storage<br>Gateway VM | Amazon Web<br>Services | Transmiss<br>ion Control<br>Protocol<br>(TCP) | 443 (HTTPS) | For<br>communica<br>tion from<br>an Storage<br>Gateway<br>VM to an<br>Amazon<br>service |

Port Requirements API Version 2013-06-30 313

| From | То | Protocol | Port | How Used            |
|------|----|----------|------|---------------------|
|      |    |          |      | endpoint. For       |
|      |    |          |      | information         |
|      |    |          |      | about service       |
|      |    |          |      | endpoints,          |
|      |    |          |      | see <u>Allowing</u> |
|      |    |          |      | Amazon              |
|      |    |          |      | Storage             |
|      |    |          |      | Gateway             |
|      |    |          |      | access              |
|      |    |          |      | through             |
|      |    |          |      | firewalls and       |
|      |    |          |      | <u>routers</u> .    |

| From                  | То                                     | Protocol                                  | Port     | How Used  |
|-----------------------|--|---|----------|---|
|                       |  |   |          | Console, the host from which you connect to the console must have access to your gateway's port 80. |
| Storage<br>Gateway VM | Domain<br>Name Service<br>(DNS) server | User<br>Datagram<br>Protocol<br>(UDP)/UDP | 53 (DNS) | For<br>communica<br>tion between<br>a Storage<br>Gateway VM<br>and the DNS<br>server.               |

| From                  | То                                 | Protocol | Port      | How Used  |
|-----------------------|------------------------------------|----------|-----------|---|
| Storage<br>Gateway VM | Network Time Protocol (NTP) server | UDP      | 123 (NTP) | Used by local systems to synchroni ze VM time to the host time. A Storage Gateway VM is configure d to use the following NTP servers:  • O.amazon. pool.ntp. org  • 1.amazon. pool.ntp. org  • 2.amazon. pool.ntp. org  • 3.amazon. pool.ntp. org  • Mote Not required for gateways hosted on |

| From  | То  | Protocol | Port        | How Used                         |
|---|---|----------|-------------|----------------------------------|
|   |   |          |             | Amazon<br>EC2.                   |
| Amazon<br>Storage<br>Gateway<br>Hardware<br>Appliance | Hypertext<br>Transfer<br>Protocol<br>(HTTP) proxy | ТСР      | 8080 (HTTP) | Required briefly for activation. |

The following table lists the required ports that must be opened for a File Gateway using either the Network File System (NFS) or Server Message Block (SMB) protocol. These port rules are part of your security group definition.

| Ru | Network<br>Element   | File<br>Share<br>Type | Protocol         | Port | Inbou | Outbo | Requir | Notes   |
|----|----------------------|-----------------------|------------------|------|-------|-------|--------|---|
| 1  | File share<br>client | NFS                   | TCP/UDP<br>Data  | 111  | ✓     | ✓     | ✓      | File sharing data<br>transfer (for NFS<br>only)   |
|    |                      |                       | TCP/UDP<br>NFS   | 2049 | ✓     | ✓     | ✓      | File sharing data<br>transfer (for NFS<br>only)   |
|    |                      |                       | TCP/UDP<br>NFSv3 | 2004 | ✓     | ✓     | ✓      | File sharing data<br>transfer (for NFS<br>only)   |
|    |                      | SMB                   | TCP/UDP<br>SMBv2 | 139  | ✓     | ✓     | ✓      | File sharing data<br>transfer session<br>service (for SMB<br>only); replaces<br>ports 137–139 for |

| Ru | Network<br>Element | File<br>Share<br>Type | Protocol         | Port | Inbou | Outbo | Requir | Notes  |
|----|--------------------|-----------------------|------------------|------|-------|-------|--------|--|
|    |                    |                       |                  |      |       |       |        | Microsoft Windows<br>NT and later  |
|    |                    |                       | TCP/UDP<br>SMBv3 | 445  | ✓     | ✓     | ✓      | File sharing data<br>transfer session<br>service (for SMB<br>only); replaces<br>ports 137–139 for<br>Microsoft Windows<br>NT and later |
| 2  | Web browser        | NFS and<br>SMB        | TCP HTTP         | 80   | ✓     | ✓     | ✓      | Amazon Web Services Management Console (activation only)   |
|    |                    |                       | TCP HTTPS        | 443  | ✓     | ✓     | ✓      | Amazon Web Services Management Console (all other operations)  |
| 3  | DNS                | NFS and<br>SMB        | TCP/UDP<br>DNS   | 53   | ✓     | ✓     | ✓      | IP name resolution   |

| Ru | Network<br>Element               | File<br>Share<br>Type | Protocol            | Port | Inbou    | Outbo    | Requir | Notes   |
|----|----------------------------------|-----------------------|---------------------|------|----------|----------|--------|---|
| 4  | NTP                              | NFS and<br>SMB        | UDP NTP             | 123  | <b>√</b> | ✓        |        | Time synchroni zation service  Note Not required for gateways hosted on Amazon EC2. |
| 5  | Microsoft<br>Active<br>Directory | SMB                   | UDP<br>NetBIOS      | 137  | ✓        | ✓        | ✓      | Name service (not used for NFS)   |
|    |                                  |                       | UDP<br>NetBIOS      | 138  | ✓        | ✓        | ✓      | Datagram service  |
|    |                                  |                       | TCP/UDP<br>LDAP     | 389  | ✓        | <b>√</b> |        | Microsoft Active Directory - Directory System Agent (DSA); client connection        |
|    |                                  |                       | TCP/UDP<br>Kerberos | 88   | ✓        | ✓        |        |   |

| Ru | Network<br>Element   | File<br>Share<br>Type | Protocol  | Port | Inbou | Outbo | Requir | Notes                    |
|----|----------------------|-----------------------|---|------|-------|-------|--------|--------------------------|
|    |                      |                       | TCP Distributed Computing Environme nt/End Point Mapper (DCE/ EMAP) | 135  | ✓     | ✓     |        |                          |
| 6  | Amazon S3            | NFS and<br>SMB        | HTTPS<br>data   | 443  | ✓     | ✓     | ✓      | Storage data<br>transfer |
| 7  | Storage<br>Gateway   | NFS and<br>SMB        | TCP SSH   | 22   | ✓     | ✓     | ✓      | Support channel          |
|    |                      |                       | TCP HTTPS   | 443  | ✓     | ✓     | ✓      | Management control       |
| 8  | Amazon<br>CloudFront | NFS and<br>SMB        | TCP HTTPS   | 443  | ✓     | ✓     | ✓      | For activation           |

## **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 gateway, make sure that you have an Activate Directory service account with delegated permissions to join computers to your domain. A service account is a user account that has been delegated certain tasks. You provide the username and password credentials for this account when you join a gateway to your Active Directory domain.

The 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

Active Directory permissions API Version 2013-06-30 322

- · 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.

## **Connecting to Your Gateway**

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: <u>Accessing the Gateway Local Console with Linux KVM</u>
- 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.amazonaws.cn/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.amazonaws.cn/ec2/.
- 2. In the navigation pane, choose **Instances**, and then select the EC2 instance that your gateway is deployed on.
- 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<br>Type  | ARN Format                 |                                      |
|-------------------|----------------------------|--------------------------------------|
| Gateway<br>ARN    | arn:aws:storagegateway: id | region:account-id :gateway/ gateway- |
| File Share<br>ARN | arn:aws:storagegateway:    | region:account-id :share/share-id    |

## **Working with Resource IDs**

When you create a resource, Storage Gateway assigns the resource a unique resource ID. This resource ID is part of the resource ARN. A resource ID takes the form of a resource identifier, followed by a hyphen, and a unique combination of eight letters and numbers. For example, a gateway ID is of the form sgw-12A3456B where sgw is the resource identifier for gateways.

Storage Gateway resource IDs are in uppercase. However, when you use these resource IDs with the Amazon EC2 API, Amazon EC2 expects resource IDs in lowercase. You must change your resource ID to lowercase to use it with the EC2 API. For example, in Storage Gateway the ID for a volume might be vol-1122AABB. When you use this ID with the EC2 API, you must change it to vol-1122aabb. Otherwise, the EC2 API might not behave as expected.

## ▲ Important

IDs for Storage Gateway volumes and Amazon EBS snapshots created from gateway volumes are changing to a longer format. Starting in December 2016, all new volumes and

snapshots will be created with a 17-character string. Starting in April 2016, you will be able to use these longer IDs so you can test your systems with the new format. For more information, see Longer EC2 and EBS Resource IDs.

For example, a volume ARN with the longer volume ID format will look like this: arn:aws:storagegateway:us-west-2:111122223333:gateway/sgw-12A3456B/volume/vol-1122AABBCCDDEEFFG.

A snapshot ID with the longer ID format will look like this: snap-78e226633445566ee. For more information, see <u>Announcement: Heads-up – Longer Storage Gateway volume</u> 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 Amazon use.

Tagging Your Resources API Version 2013-06-30 326

 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

- Open the Storage Gateway console at https://console.amazonaws.cn/storagegateway/home.
- In the navigation pane, choose the resource you want to tag. 2.

For example, to tag a gateway, choose **Gateways**, and then choose the gateway you want to tag from the list of gateways.

- Choose **Tags**, and then choose **Add/edit tags**. 3.
- In the **Add/edit tags** dialog box, choose **Create tag**. 4.
- 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

- Open the Storage Gateway console at https://console.amazonaws.cn/storagegateway/home. 1.
- 2. Choose the resource whose tag you want to edit.
- 3. Choose **Tags** to open the **Add/edit tags** dialog box.
- Choose the pencil icon next to the tag you want to edit, and then edit the tag. 4.
- When you're done editing the tag, choose **Save**. 5.

Working with tags API Version 2013-06-30 327

#### To delete a tag

- 1. Open the Storage Gateway console at <a href="https://console.amazonaws.cn/storagegateway/home">https://console.amazonaws.cn/storagegateway/home</a>.
- 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**.

#### See also

Using tags to control access to your gateway and resources

# Working with open-source components for Amazon Storage Gateway

This section describes third-party tools and licenses that we depend on to deliver 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 Amazon Storage Gateway software:

- For gateways deployed on VMware ESXi: <u>sources.tar</u>
- For gateways deployed on Microsoft Hyper-V: sources\_hyperv.tar
- For gateways deployed on Linux Kernel-based Virtual Machine (KVM): sources\_KVM.tar

See also API Version 2013-06-30 328

This product includes software developed by the OpenSSL project for use in the OpenSSL Toolkit (<a href="http://www.openssl.org/">http://www.openssl.org/</a>). For the relevant licenses for all dependent third-party tools, see <a href="http://www.openssl.org/">Third-Party Licenses</a>.

## 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 (<a href="http://www.openssl.org/">http://www.openssl.org/</a>). For the relevant licenses for all dependent third-party tools, see <a href="http://www.openssl.org/">Third-Party Licenses</a>.

## Quotas

## **Quotas for file shares**

The following table lists quotas for file shares.

| Limit |
|-------|
| 50    |
|       |
|       |

# **Description** Limit The number of file shares managed by a gateway can impact the gateway's performance. For more informati on, see Performance guidance for gateways with multiple file shares. The maximum size of an individual file, which 5 TB is the maximum size of an individual object in Amazon S3 Note If you write a file larger than 5 TB, you get a "file too large" error message and only the first 5 TB of the file is uploaded. Maximum path length 1024 bytes Note 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.

## 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<br>Local Disks |
|-----------------|-----------------|-----------------|-------------------------------|
| S3 File Gateway | 150 GiB         | 64 TiB          | _                             |

#### (i) 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.

Using storage classes API Version 2013-06-30 331

| Amazon S3 storage class | Considerations  |
|-------------------------|---|
| Standard                | Choose Standard to store your frequently accessed files redundantly in multiple Availabil ity Zones that are geographically separated. This is the default storage class. See Amazon S3 pricing for more details.   |
| S3 Intelligent-Tiering  | Choose Intelligent-Tiering to optimize storage costs by automatically moving data to the most cost-effective storage access tier.  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.  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. |

| Amazon S3 storage class | Considerations   |
|-------------------------|--|
|                         | 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 operations impact object creation before using |
|                         | this storage class in production. See Amazon S3 pricing for more details.  |

| Amazon S3 storage class | Considerations  |
|-------------------------|---|
| S3 Standard-IA          | Choose Standard-IA to store your infrequently accessed files redundantly in multiple Availabil ity Zones that are geographically separated.   |
|                         | 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. |
|                         | 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.   |

| Amazon S3 storage class | Considerations   |
|-------------------------|--|
| S3 One Zone-IA          | Choose One Zone-IA to store your infrequently accessed files in a single Availability Zone.  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. |
|                         | 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.  |

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 <a href="Object lifecycle management">Object lifecycle management</a>.

## 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 uses a different cost structure than basic Amazon S3 storage, 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. 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 an object or objects. Depending on the type of provisioning that you want Kubernetes to use when your pods request storage, 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**

- SMB CSI drivers
- NFS CSI drivers

#### SMB CSI drivers

Follow the procedures in this section to install and configure 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:

- 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 --
- 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:

SMB CSI drivers API Version 2013-06-30 337

| NAME                                |                          | READY | STATUS  | RESTARTS | AGE | IP |
|-------------------------------------|--------------------------|-------|---------|----------|-----|----|
| NOI                                 | DE                       |       |         |          |     |    |
| csi-smb-contro                      | ller-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: "examplenamespace"
    csi.storage.k8s.io/node-stage-secret-name: "examplesmbcreds"
    csi.storage.k8s.io/node-stage-secret-namespace: "examplenamespace"
volumeBindingMode: Immediate
reclaimPolicy: Retain
mountOptions:
    - dir_mode=0777
    - file_mode=0777
    - uid=1001
```

SMB CSI drivers API Version 2013-06-30 338

- gid=1001
- From a command-line terminal with access to kubectl and storageclass.yaml, run the 2. 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.

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:

- Create two configuration files. One named persistent volume. yaml, and one named persistentvolumeclaim.yaml.
- For persistent volume. 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
```

SMB CSI drivers API Version 2013-06-30 339

```
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
```

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 specfied volumeName matches the
name of the PersistentVolume you created
        storageClassName: ""
```

4. From a command-line terminal with access to kubect1 and both of the .yaml files that you created, run the following commands:

kubectl apply -f persistentvolume.yaml

kubectl apply -f persistentvolumeclaim.yaml

SMB CSI drivers API Version 2013-06-30 340



#### 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.

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 -skSL https://raw.githubusercontent.com/kubernetes-csi/csi-driver-smb/master/ deploy/uninstall-driver.sh | bash -s --

## NFS CSI drivers

Follow the procedures in this section to install and configure 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:

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 --

NFS CSI drivers API Version 2013-06-30 341 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 | ΙP |
|-------------------------------------|--------------------------|-------|---------|----------|-----|----|
| NO                                  | DE                       |       |         |          |     |    |
| csi-nfs-contro                      | ller-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-cvgbs                  |                          | 3/3   | Running | 0        | 35s |    |
| 10.240.0.35                         | k8s-agentpool-22533604-1 |       |         |          |     |    |
| csi-nfs-node-d                      | r4s4                     | 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
```

NFS CSI drivers API Version 2013-06-30 342

```
volumeBindingMode: Immediate
mountOptions:
    - hard
    - nfsvers=4.1
```

From a command-line terminal with access to kubectl and storageclass.yaml, run the 2. 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.

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:

- Create two configuration files named persistent volume. yaml and persistentvolumeclaim.yaml.
- For persistent volume. 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-nfs-examplename
spec:
    capacity:
        storage: 10Gi
    accessModes:
        - ReadWriteMany
```

NFS CSI drivers API Version 2013-06-30 343

```
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: examplename-pvc-nfs-static
spec:
   accessModes:
        - ReadWriteMany
   resources:
        requests:
        storage: 10Gi
   volumeName: pv-nfs-examplename # make sure specfied 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

NFS CSI drivers API Version 2013-06-30 344



#### 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.

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 --

## **Amazon 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.

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

Terraform module API Version 2013-06-30 345 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 Amazon service. Best-effort support is provided by the Amazon Storage community.

API Version 2013-06-30 346 Terraform module

## **API Reference for Storage Gateway**

In addition to using the console, you can use the Amazon Storage Gateway API to programmatically configure and manage your gateways. This section describes the Amazon Storage Gateway operations, request signing for authentication and the error handling. For information about the regions and endpoints available for Storage Gateway, see Amazon Storage Gateway Endpoints and Quotas in the Amazon Web Services General Reference.

#### Note

You can also use the Amazon SDKs when developing applications with Storage Gateway. The Amazon 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**

- Amazon Storage Gateway Required Request Headers
- Signing Requests
- **Error Responses**
- Storage Gateway API Actions

## **Amazon Storage Gateway Required Request Headers**

This section describes the required headers that you must send with every POST request to Amazon 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.cn

Content-Type: application/x-amz-json-1.1

Required Request Headers API Version 2013-06-30 347 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 Amazon Storage Gateway. Headers shown below that begin with "x-amz" are Amazon-specific headers. All other headers listed are common header used in HTTP transactions.

| Header        | Description   |
|---------------|---|
| Authorization | The authorization header contains several of pieces of information about the request that allow Amazon Storage Gateway to determine if the request is a valid action for the requester. The format of this header is as follows (line breaks added for readability):  |
|               | Authorization: AWS4-HMAC_SHA456 Credentials= YourAccessKey /yyymmdd/region/storagegateway/aws4_request, SignedHeaders=content-type;host;x-amz-date;x-amz-target, Signature= CalculatedSignature   |
|               | In the preceding syntax, you specify <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 Amazon V4 Signing process. The details of signing are discussed in the topic <u>Signing Requests</u> . |
| Content-Type  | Use application/x-amz-json-1.1 as the content type for all requests to Amazon Storage Gateway.  |
|               | Content-Type: application/x-amz-json-1.1  |
| Host          | Use the host header to specify the Amazon Storage Gateway endpoint where you send your request. For example, storagegateway.us-   |

Required Request Headers API Version 2013-06-30 348

| Header       | Description  |
|--------------|--|
|              | east-2.amazonaws.com is the endpoint for the US East (Ohio) region. For more information about the endpoints available for Amazon Storage Gateway, see <a href="Amazon Storage Gateway Endpoints">Amazon Storage Gateway Endpoints</a> and Quotas in the Amazon Web Services General Reference.  Host: storagegateway. region.amazonaws.com  |
| x-amz-date   | You must provide the time stamp in either the HTTP Date header or the AWS x-amz-date header. (Some HTTP client libraries don't let you set the Date header.) When an x-amz-date header is present, the Amazon Storage Gateway ignores any Date header during the request authentication. The x-amz-date format must be ISO8601 Basic in the YYYYMMDD'T'HHMMSS'Z' format. If both the Date and x-amz-date header are used, the format of the Date header does not have to be ISO8601. |
|              | x-amz-date: YYYYMMDD'T'HHMMSS'Z'   |
| x-amz-target | This header specifies the version of the API and the operation that you are requesting. The target header values are formed by concatenating the API version with the API name and are in the following format.  |
|              | x-amz-target: StorageGateway_ APIversion .operationName  |
|              | The <i>operationName</i> value (e.g. "ActivateGateway") can be found from the API list, <u>API Reference for Storage Gateway</u> .   |

# **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

Signing Requests API Version 2013-06-30 349

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 <u>Amazon Signature Version 4</u>. 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 <u>ListGateways</u>. 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: 20120910T0000000Z
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:20120910T0000000Z
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 <u>Task 2</u>: <u>Create a String to Sign</u> is:

```
AWS4-HMAC-SHA256
20120910T000000Z
20120910/us-east-2/storagegateway/aws4_request
92c0effa6f9224ac752ca179a04cecbede3038b0959666a8160ab452c9e51b3e
```

The first line of the *string to sign* is the algorithm, the second line is the time stamp, the third line is the *credential scope*, and the last line is a hash of the canonical request from Task 1.

For Task 3: Create a Signature, the *derived key* can be represented as:

```
derived key = HMAC(HMAC(HMAC("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 Amazon 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.

Error Responses API Version 2013-06-30 352

### **Exceptions**

The following table lists Amazon Storage Gateway API exceptions. When an Amazon 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             |
|--|--|------------------------------|
| <pre>IncompleteSignatur eException</pre>   | The specified signature is incomplete.   | 400 Bad Request              |
| InternalFailure                            | The request processing has failed due to some unknown error, exception or failure.   | 500 Internal<br>Server Error |
| InternalServerError                        | One of the operation error code messages Operation Error Codes.  | 500 Internal<br>Server Error |
| InvalidAction                              | The requested action or operation is invalid.  | 400 Bad Request              |
| InvalidClientTokenId                       | The X.509 certificate or Amazon Access Key ID provided does not exist in our records.  | 403 Forbidden                |
| <pre>InvalidGatewayRequ estException</pre> | One of the operation error code messages in <u>Operation Error Codes</u> .   | 400 Bad Request              |
| InvalidSignatureEx ception                 | The request signature we calculate d does not match the signature you provided. Check your Amazon Access Key and signing method. | 400 Bad Request              |
| MissingAction                              | The request is missing an action or operation parameter.   | 400 Bad Request              |

Exceptions API Version 2013-06-30 353

| Exception                       | Message   | HTTP Status Code           |
|---------------------------------|---|----------------------------|
| MissingAuthenticat<br>ionToken  | The request must contain either a valid (registered) Amazon 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 serializa tion. 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<br>Unavailable |
| SubscriptionRequir edException  | The Amazon Access Key Id needs a subscription for the service.  | 400 Bad Request            |
| ThrottlingException             | Rate exceeded.  | 400 Bad Request            |
| UnknownOperationEx ception      | An unknown operation was specified. Valid operations are listed in <a href="Storage">Storage</a> Gateway API Actions.                                       | 400 Bad Request            |
| UnrecognizedClient<br>Exception | 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 Amazon Storage Gateway operation error codes and APIs that can return the codes. All operation error codes are returned with one of two general

 ${\bf exceptions-} Internal Server Error \ and \ Invalid Gateway Request Exception-described \ in \\ {\bf Exceptions.}$ 

| Operation Error Code                  | Message   | Operations That Return this<br>Error Code                     |
|---------------------------------------|---|---|
| ActivationKeyExpired                  | The specified activation key has expired.       | <u>ActivateGateway</u>  |
| ActivationKeyInvalid                  | The specified activation hey is invalid.        | <u>ActivateGateway</u>  |
| ActivationKeyNotFound                 | The specified activation key was not found.     | <u>ActivateGateway</u>  |
| BandwidthThrottleS<br>cheduleNotFound | The specified bandwidth throttle was not found. | DeleteBandwidthRateLimit                                      |
| CannotExportSnapshot                  | The specified snapshot cannot be exported.      | <u>CreateCachediSCSIVolume</u> <u>CreateStorediSCSIVolume</u> |
| InitiatorNotFound                     | The specified initiator was not found.          | <u>DeleteChapCredentials</u>                                  |
| DiskAlreadyAllocated                  | The specified disk is already allocated.        | AddUploadBuffer  AddWorkingStorage  CreateStorediSCSIVolume   |
| DiskDoesNotExist                      | The specified disk does not exist.              | AddUploadBuffer  AddWorkingStorage  CreateStorediSCSIVolume   |

| Operation Error Code                                   | Message   | Operations That Return this<br>Error Code   |
|--|---|---|
| DiskSizeNotGigAligned                                  | The specified disk is not gigabyte-aligned.   | CreateStorediSCSIVolume   |
| DiskSizeGreaterTha<br>nVolumeMaxSize                   | The specified disk size is greater than the maximum volume size.                                | CreateStorediSCSIVolume   |
| DiskSizeLessThanVo<br>lumeSize                         | The specified disk size is less than the volume size.   | CreateStorediSCSIVolume   |
| DuplicateCertifica<br>teInfo                           | The specified certifica te information is a duplicate.  | <u>ActivateGateway</u>  |
| FileSystemAssociationEndpoi<br>ntConfigurationConflict | Existing File System Association endpoint configuration conflicts with specified configuration. | <u>AssociateFileSystem</u>  |
| FileSystemAssociationEndpoi<br>ntIpAddressAlreadyInUse | The specified endpoint IP address is already in use.  | <u>AssociateFileSystem</u>  |
| FileSystemAssociationEndpoi<br>ntIpAddressMissing      | File System Associati on Endpoint IP address is missing.  | <u>AssociateFileSystem</u>  |
| FileSystemAssociationNotFound                          | The specified file system association was not found.  | UpdateFileSystemAssociation  DisassociateFileSystem  DescribeFileSystemAssociations |

| Operation Error Code | Message                                  | Operations That Return this<br>Error Code |
|----------------------|--|---|
| FileSystemNotFound   | The specified file system was not found. | <u>AssociateFileSystem</u>                |

| Operation Error Code | Message                            | Operations That Return this<br>Error Code    |
|----------------------|------------------------------------|--|
| GatewayInternalError | A gateway internal error occurred. | AddCache                                     |
|                      |                                    | <u>AddUploadBuffer</u>                       |
|                      |                                    | <u>AddWorkingStorage</u>                     |
|                      |                                    | CreateCachediSCSIVolume                      |
|                      |                                    | CreateSnapshot                               |
|                      |                                    | CreateStorediSCSIVolume                      |
|                      |                                    | <u>CreateSnapshotFromVolumeRecoveryPoint</u> |
|                      |                                    | DeleteBandwidthRateLimit                     |
|                      |                                    | DeleteChapCredentials                        |
|                      |                                    | <u>DeleteVolume</u>                          |
|                      |                                    | DescribeBandwidthRateLimit                   |
|                      |                                    | <u>DescribeCache</u>                         |
|                      |                                    | DescribeCachediSCSIVolumes                   |
|                      |                                    | <u>DescribeChapCredentials</u>               |
|                      |                                    | DescribeGatewayInformation                   |
|                      |                                    | <u>DescribeMaintenanceStartTime</u>          |
|                      |                                    | <u>DescribeSnapshotSchedule</u>              |
|                      |                                    | DescribeStorediSCSIVolumes                   |
|                      |                                    | <u>DescribeWorkingStorage</u>                |
|                      |                                    | ListLocalDisks                               |

| Operation Error Code | Message | Operations That Return this<br>Error Code |
|----------------------|---------|---|
|                      |         | ListVolumes                               |
|                      |         | ListVolumeRecoveryPoints                  |
|                      |         | ShutdownGateway                           |
|                      |         | StartGateway                              |
|                      |         | <u>UpdateBandwidthRateLimit</u>           |
|                      |         | <u>UpdateChapCredentials</u>              |
|                      |         | <u>UpdateMaintenanceStartTime</u>         |
|                      |         | <u>UpdateGatewaySoftwareNow</u>           |
|                      |         | <u>UpdateSnapshotSchedule</u>             |

| Operation Error Code | Message                                 | Operations That Return this<br>Error Code    |
|----------------------|---|--|
| GatewayNotConnected  | The specified gateway is not connected. | AddCache                                     |
|                      |   | <u>AddUploadBuffer</u>                       |
|                      |   | <u>AddWorkingStorage</u>                     |
|                      |   | CreateCachediSCSIVolume                      |
|                      |   | CreateSnapshot                               |
|                      |   | CreateStorediSCSIVolume                      |
|                      |   | <u>CreateSnapshotFromVolumeRecoveryPoint</u> |
|                      |   | DeleteBandwidthRateLimit                     |
|                      |   | DeleteChapCredentials                        |
|                      |   | <u>DeleteVolume</u>                          |
|                      |   | DescribeBandwidthRateLimit                   |
|                      |   | <u>DescribeCache</u>                         |
|                      |   | <u>DescribeCachediSCSIVolumes</u>            |
|                      |   | <u>DescribeChapCredentials</u>               |
|                      |   | DescribeGatewayInformation                   |
|                      |   | <u>DescribeMaintenanceStartTime</u>          |
|                      |   | <u>DescribeSnapshotSchedule</u>              |
|                      |   | DescribeStorediSCSIVolumes                   |
|                      |   | <u>DescribeWorkingStorage</u>                |
|                      |   | ListLocalDisks                               |

| Operation Error Code | Message | Operations That Return this<br>Error Code |
|----------------------|---------|---|
|                      |         | ListVolumes                               |
|                      |         | ListVolumeRecoveryPoints                  |
|                      |         | ShutdownGateway                           |
|                      |         | <u>StartGateway</u>                       |
|                      |         | <u>UpdateBandwidthRateLimit</u>           |
|                      |         | <u>UpdateChapCredentials</u>              |
|                      |         | <u>UpdateMaintenanceStartTime</u>         |
|                      |         | UpdateGatewaySoftwareNow                  |
|                      |         | <u>UpdateSnapshotSchedule</u>             |

| Operation Error Code | Message                              | Operations That Return this<br>Error Code               |
|----------------------|--------------------------------------|---|
| GatewayNotFound      | The specified gateway was not found. | AddCache  |
|                      |                                      | <u>AddUploadBuffer</u>                                  |
|                      |                                      | <u>AddWorkingStorage</u>                                |
|                      |                                      | CreateCachediSCSIVolume                                 |
|                      |                                      | CreateSnapshot  |
|                      |                                      | <u>CreateSnapshotFromVolumeRec</u><br><u>overyPoint</u> |
|                      |                                      | CreateStorediSCSIVolume                                 |
|                      |                                      | DeleteBandwidthRateLimit                                |
|                      |                                      | DeleteChapCredentials                                   |
|                      |                                      | DeleteGateway   |
|                      |                                      | <u>DeleteVolume</u>                                     |
|                      |                                      | DescribeBandwidthRateLimit                              |
|                      |                                      | <u>DescribeCache</u>                                    |
|                      |                                      | <u>DescribeCachediSCSIVolumes</u>                       |
|                      |                                      | <u>DescribeChapCredentials</u>                          |
|                      |                                      | DescribeGatewayInformation                              |
|                      |                                      | <u>DescribeMaintenanceStartTime</u>                     |
|                      |                                      | <u>DescribeSnapshotSchedule</u>                         |
|                      |                                      | DescribeStorediSCSIVolumes                              |
|                      |                                      | <u>DescribeWorkingStorage</u>                           |

| Operation Error Code | Message | Operations That Return this<br>Error Code |
|----------------------|---------|---|
|                      |         | <u>ListLocalDisks</u>                     |
|                      |         | ListVolumes                               |
|                      |         | ListVolumeRecoveryPoints                  |
|                      |         | ShutdownGateway                           |
|                      |         | <u>StartGateway</u>                       |
|                      |         | <u>UpdateBandwidthRateLimit</u>           |
|                      |         | <u>UpdateChapCredentials</u>              |
|                      |         | <u>UpdateMaintenanceStartTime</u>         |
|                      |         | <u>UpdateGatewaySoftwareNow</u>           |
|                      |         | <u>UpdateSnapshotSchedule</u>             |

| Operation Error Code | Message                           | Operations That Return this<br>Error Code    |
|----------------------|-----------------------------------|--|
| GatewayProxyNetwor   | The specified gateway             | AddCache                                     |
| kConnectionBusy      | proxy network connection is busy. | <u>AddUploadBuffer</u>                       |
|                      |                                   | AddWorkingStorage                            |
|                      |                                   | CreateCachediSCSIVolume                      |
|                      |                                   | CreateSnapshot                               |
|                      |                                   | <u>CreateSnapshotFromVolumeRecoveryPoint</u> |
|                      |                                   | CreateStorediSCSIVolume                      |
|                      |                                   | DeleteBandwidthRateLimit                     |
|                      |                                   | DeleteChapCredentials                        |
|                      |                                   | <u>DeleteVolume</u>                          |
|                      |                                   | DescribeBandwidthRateLimit                   |
|                      |                                   | <u>DescribeCache</u>                         |
|                      |                                   | DescribeCachediSCSIVolumes                   |
|                      |                                   | <u>DescribeChapCredentials</u>               |
|                      |                                   | DescribeGatewayInformation                   |
|                      |                                   | DescribeMaintenanceStartTime                 |
|                      |                                   | <u>DescribeSnapshotSchedule</u>              |
|                      |                                   | DescribeStorediSCSIVolumes                   |
|                      |                                   | <u>DescribeWorkingStorage</u>                |
|                      |                                   | ListLocalDisks                               |

| Operation Error Code | Message | Operations That Return this<br>Error Code |  |
|----------------------|---------|---|--|
|                      |         | ListVolumes                               |  |
|                      |         | ListVolumeRecoveryPoints                  |  |
|                      |         | ShutdownGateway                           |  |
|                      |         | <u>StartGateway</u>                       |  |
|                      |         | <u>UpdateBandwidthRateLimit</u>           |  |
|                      |         | <u>UpdateChapCredentials</u>              |  |
|                      |         | <u>UpdateMaintenanceStartTime</u>         |  |
|                      |         | UpdateGatewaySoftwareNow                  |  |
|                      |         | <u>UpdateSnapshotSchedule</u>             |  |

| Operation Error Code | Message           | Operations That Return this<br>Error Code               |
|----------------------|-------------------|---|
| InternalError        | An internal error | ActivateGateway   |
|                      | occurred.         | AddCache  |
|                      |                   | <u>AddUploadBuffer</u>                                  |
|                      |                   | <u>AddWorkingStorage</u>                                |
|                      |                   | CreateCachediSCSIVolume                                 |
|                      |                   | CreateSnapshot  |
|                      |                   | <u>CreateSnapshotFromVolumeRec</u><br><u>overyPoint</u> |
|                      |                   | CreateStorediSCSIVolume                                 |
|                      |                   | <u>DeleteBandwidthRateLimit</u>                         |
|                      |                   | <u>DeleteChapCredentials</u>                            |
|                      |                   | <u>DeleteGateway</u>                                    |
|                      |                   | <u>DeleteVolume</u>                                     |
|                      |                   | DescribeBandwidthRateLimit                              |
|                      |                   | <u>DescribeCache</u>                                    |
|                      |                   | DescribeCachediSCSIVolumes                              |
|                      |                   | DescribeChapCredentials                                 |
|                      |                   | DescribeGatewayInformation                              |
|                      |                   | <u>DescribeMaintenanceStartTime</u>                     |
|                      |                   | <u>DescribeSnapshotSchedule</u>                         |
|                      |                   | <u>DescribeStorediSCSIVolumes</u>                       |

| Operation Error Code | Message | Operations That Return this<br>Error Code |
|----------------------|---------|---|
|                      |         | <u>DescribeWorkingStorage</u>             |
|                      |         | ListLocalDisks                            |
|                      |         | ListGateways                              |
|                      |         | ListVolumes                               |
|                      |         | ListVolumeRecoveryPoints                  |
|                      |         | ShutdownGateway                           |
|                      |         | <u>StartGateway</u>                       |
|                      |         | <u>UpdateBandwidthRateLimit</u>           |
|                      |         | <u>UpdateChapCredentials</u>              |
|                      |         | <u>UpdateMaintenanceStartTime</u>         |
|                      |         | <u>UpdateGatewayInformation</u>           |
|                      |         | <u>UpdateGatewaySoftwareNow</u>           |
|                      |         | <u>UpdateSnapshotSchedule</u>             |

| Operation Error Code | Message                      | Operations That Return this<br>Error Code    |
|----------------------|------------------------------|--|
| InvalidParameters    | The specified request        | ActivateGateway                              |
|                      | contains invalid parameters. | AddCache                                     |
|                      |                              | AddUploadBuffer                              |
|                      |                              | <u>AddWorkingStorage</u>                     |
|                      |                              | CreateCachediSCSIVolume                      |
|                      |                              | CreateSnapshot                               |
|                      |                              | <u>CreateSnapshotFromVolumeRecoveryPoint</u> |
|                      |                              | CreateStorediSCSIVolume                      |
|                      |                              | DeleteBandwidthRateLimit                     |
|                      |                              | DeleteChapCredentials                        |
|                      |                              | <u>DeleteGateway</u>                         |
|                      |                              | <u>DeleteVolume</u>                          |
|                      |                              | DescribeBandwidthRateLimit                   |
|                      |                              | DescribeCache                                |
|                      |                              | DescribeCachediSCSIVolumes                   |
|                      |                              | DescribeChapCredentials                      |
|                      |                              | DescribeGatewayInformation                   |
|                      |                              | DescribeMaintenanceStartTime                 |
|                      |                              | <u>DescribeSnapshotSchedule</u>              |
|                      |                              | DescribeStorediSCSIVolumes                   |

| Operation Error Code | Message                       | Operations That Return this<br>Error Code |
|----------------------|-------------------------------|---|
|                      |                               | <u>DescribeWorkingStorage</u>             |
|                      |                               | ListLocalDisks                            |
|                      |                               | ListGateways                              |
|                      |                               | ListVolumes                               |
|                      |                               | ListVolumeRecoveryPoints                  |
|                      |                               | ShutdownGateway                           |
|                      |                               | <u>StartGateway</u>                       |
|                      |                               | <u>UpdateBandwidthRateLimit</u>           |
|                      |                               | <u>UpdateChapCredentials</u>              |
|                      |                               | <u>UpdateMaintenanceStartTime</u>         |
|                      |                               | <u>UpdateGatewayInformation</u>           |
|                      |                               | <u>UpdateGatewaySoftwareNow</u>           |
|                      |                               | <u>UpdateSnapshotSchedule</u>             |
| LocalStorageLimitE   | The local storage limit       | AddCache                                  |
| xceeded              | was exceeded.                 | AddUploadBuffer                           |
|                      |                               | AddWorkingStorage                         |
| LunInvalid           | The specified LUN is invalid. | CreateStorediSCSIVolume                   |

| Operation Error Code            | Message  | Operations That Return this<br>Error Code   |
|---------------------------------|--|---|
| MaximumVolumeCount<br>Exceeded  | The maximum volume count was exceeded.         | CreateCachediSCSIVolume CreateStorediSCSIVolume DescribeCachediSCSIVolumes DescribeStorediSCSIVolumes |
| NetworkConfigurati<br>onChanged | The gateway network configuration has changed. | <u>CreateCachediSCSIVolume</u> <u>CreateStorediSCSIVolume</u>   |

| Operation Error Code | Message                     | Operations That Return this<br>Error Code    |
|----------------------|-----------------------------|--|
| NotSupported         | The specified               | ActivateGateway                              |
|                      | operation is not supported. | AddCache                                     |
|                      |                             | <u>AddUploadBuffer</u>                       |
|                      |                             | AddWorkingStorage                            |
|                      |                             | CreateCachediSCSIVolume                      |
|                      |                             | CreateSnapshot                               |
|                      |                             | <u>CreateSnapshotFromVolumeRecoveryPoint</u> |
|                      |                             | CreateStorediSCSIVolume                      |
|                      |                             | DeleteBandwidthRateLimit                     |
|                      |                             | <u>DeleteChapCredentials</u>                 |
|                      |                             | <u>DeleteGateway</u>                         |
|                      |                             | <u>DeleteVolume</u>                          |
|                      |                             | DescribeBandwidthRateLimit                   |
|                      |                             | <u>DescribeCache</u>                         |
|                      |                             | <u>DescribeCachediSCSIVolumes</u>            |
|                      |                             | <u>DescribeChapCredentials</u>               |
|                      |                             | DescribeGatewayInformation                   |
|                      |                             | <u>DescribeMaintenanceStartTime</u>          |
|                      |                             | <u>DescribeSnapshotSchedule</u>              |
|                      |                             | DescribeStorediSCSIVolumes                   |

| Operation Error Code            | Message                                | Operations That Return this<br>Error Code   |
|---------------------------------|--|---|
|                                 |  | DescribeWorkingStorage  ListLocalDisks  ListGateways  ListVolumes  ListVolumeRecoveryPoints  ShutdownGateway  StartGateway  UpdateBandwidthRateLimit  UpdateChapCredentials  UpdateMaintenanceStartTime |
|                                 |  | UpdateGatewayInformation UpdateGatewaySoftwareNow UpdateSnapshotSchedule  |
| OutdatedGateway                 | The specified gateway is out of date.  | ActivateGateway   |
| SnapshotInProgress<br>Exception | The specified snapshot is in progress. | DeleteVolume  |
| SnapshotIdInvalid               | The specified snapshot is invalid.     | <u>CreateCachediSCSIVolume</u> <u>CreateStorediSCSIVolume</u>   |
| StagingAreaFull                 | The staging area is full.              | <u>CreateCachediSCSIVolume</u> <u>CreateStorediSCSIVolume</u>   |

| Operation Error Code | Message                              | Operations That Return this<br>Error Code  |
|----------------------|--------------------------------------|--|
| TargetAlreadyExists  | The specified target already exists. | <u>CreateCachediSCSIVolume</u> <u>CreateStorediSCSIVolume</u>  |
| 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<br>Error Code   |
|--|---|---|
| UnsupportedOperati<br>onForGatewayType | The specified operation is not valid for the type of the gateway. | AddCache  AddWorkingStorage  CreateCachediSCSIVolume  CreateSnapshotFromVolumeRecoveryPoint  CreateStorediSCSIVolume  DeleteSnapshotSchedule  DescribeCache  DescribeCachediSCSIVolumes  DescribeStorediSCSIVolumes  DescribeStorediSCSIVolumes  DescribeUploadBuffer  DescribeWorkingStorage |
| VolumeAlreadyExists                    | The specified volume already exists.                              | <u>CreateCachediSCSIVolume</u> <u>CreateStorediSCSIVolume</u>   |
| VolumeIdInvalid                        | The specified volume is invalid.                                  | DeleteVolume  |
| VolumeInUse                            | The specified volume is already in use.                           | <u>DeleteVolume</u>   |

| Operation Error Code | Message                             | Operations That Return this<br>Error Code  |
|----------------------|-------------------------------------|--|
| VolumeNotFound       | The specified volume was not found. | CreateSnapshot  CreateSnapshotFromVolumeRec overyPoint  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.

Error Responses API Version 2013-06-30 375

```
"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",
```

Error Responses API Version 2013-06-30 376

```
"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 Amazon Storage Gateway API Reference.

Actions API Version 2013-06-30 377

## Document history for the Amazon S3 File Gateway User Guide

• API version: 2013-06-30

• Latest documentation update: October 12, 2021

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            |
|--|---|-----------------|
| 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 u sing the Amazon CLI. For more information see Release Notes — Gateway Appliance Software. | October 5, 2023 |
| Updated recommended CloudWatch alarms  | The CloudWatch HealthNot ifications alarm now applies to and is recommende d for all gateway types and host platforms. Recommend ed configuration settings have also been updated for HealthNotifications and AvailabilityNotifi  | October 2, 2023 |

cations . For more information see <u>Understan</u> ding CloudWatch alarms.

### Added GatewayClockOutOfS ync 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 Amazon Storage Gateway server time. For more information, see <a href="Error: GatewayClockOutOfSync.">Error: GatewayClockOutOfSync.</a>

October 19, 2022

### <u>Updated gateway creation</u> procedures

The procedure for creating a new gateway has been updated to reflect changes in the Storage Gateway console. For more informati on, see <a href="Create and activate">Create and activate</a> 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 Managemen t Console snap-in to forceclose 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 patt erns are identified. For more information, see <a href="Understanding File Gateway">Understanding File Gateway audit logs</a>.

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 <a href="Create a file">Create a file</a> 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 Amazon 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

### Schedule-based bandwidth throttling

Storage Gateway now supports schedule-based bandwidth throttling for tape and Volume Gateways. For more information, see Scheduling bandwidth throttling using the Storage Gateway console.

November 9, 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 enumerati on of files and folders on an SMB file share based on the share's ACLs. For more information, see <u>Creating an SMB file share</u>.

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 informati on, see Replacing a File Gateway with a new File Gateway.

October 30, 2020

### File Gateway cold cache read performance 4x increase

Storage Gateway has increased cold cache read performance 4x. For more information, see <a href="Performance">Performance</a> guidance for File Gateways.

August 31, 2020

### Order the hardware appliance through the console

You can now order the hardware appliance through the Amazon Storage Gateway console. For more informati on, see <u>Using the Amazon</u>
<u>Storage Gateway Hardware</u>
Appliance.

August 12, 2020

# Support for Federal Information Processing Standard (FIPS) endpoints in new Amazon Regions

You can now activate a gateway with FIPS endpoints in the US East (Ohio), US E ast (N. Virginia), US West (N. California), US West (Oregon), and Canada (Central) Regions. For more informati on, see Amazon Storage Gateway endpoints and quotas in the Amazon Web Services 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 CloudWatc h alarms in the Storage Gateway console You can now view CloudWatch alarms in the Storage Gateway console. For more information, see <u>Understanding CloudWatch alarms</u>.

May 29, 2020

Support for Federal Informati on Processing Standard (FIPS) endpoints You can now activate a gateway with FIPS endpoints in the Amazon GovCloud (US) Regions. To choose a FIPS endpoint for a File Gateway, see Choosing a service endpoint. To choose a FIPS endpoint for a Volume Gateway, see Choosing a service endpoint. To choose a FIPS endpoint for a Tape Gateway, see Choosing a service endpoint.

May 22, 2020

**New Amazon Regions** 

Storage Gateway is now available in the Africa (Cape Town) and Europe (Milan) Regions. For more informati on, see Amazon Storage Gateway endpoints and quotas in the Amazon Web Services General Reference.

May 7, 2020

### Support for S3 Intelligent-Tiering storage class

Storage Gateway now supports S3 Intelligent-Tierin g storage class. The S3 I ntelligent-Tiering storage class optimizes storage costs by automatically moving data to the most cost-effe ctive storage access tier, without performance impact or operational overhead. For more information, see Storage class for automatic ally optimizing frequently and infrequently accessed objects in the Amazon Simple Storage Service User Guide.

April 30, 2020

#### **New Amazon Region**

Storage Gateway is now available in the Amazon GovCloud (US-East) Region. For more information, see Amazon Storage Gateway Endpoints and Quotas in the Amazon Web Services General Reference.

March 12, 2020

#### Support for Linux Kernel-ba sed 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 informat ion, see <a href="Supported Hy">Supported Hy</a> pervisors and Host Requireme <a href="mailto:nts">nts</a> 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 informat ion, 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

## New Amazon Web Services Region for Tape Gateway

Tape Gateway is now available in the South America (São Paulo) Region. For more information, see Amazon Storage Gateway Endpoints and Quotas in the Amazon Web Services General Reference.

September 24, 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 Amazon Web Services Region

Storage Gateway is now available in the Asia Pacific (Hong Kong) Region. For more information, see Amazon
Storage Gateway Endp
oints and Quotas in the
Amazon Web Services General
Reference.

August 14, 2019

#### New Amazon Web Services Region

Storage Gateway is now available in the Middle East (Bahrain) Region. For more information, see Amazon

Storage Gateway Endp
oints and Quotas in the Amazon Web Services 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 V irtual 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 <u>Using Microsoft Windows ACLs to Control Access to an SMB File Share</u>.

May 8, 2019

File Gateway support for tagbased 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 informati on, see Controlling Access to File Gateway Resources.

March 4, 2019

Availability of Amazon

Storage Gateway Hardware

Appliance in Europe

The Amazon Storage Gateway Hardware Appliance is now available in Europe. For more information, see Amazon Storage Gateway Hardware Appliance Regions in the Amazon Web Services General Reference. In addition, you can now increase the useable storage on the Amazon Storage Gateway Hardware Appliance from 5 TB to 12 TB and replace the installed cop per network card with a 10gigabit fiber optic network card. For more informat ion, see Setting Up Your Hardware Appliance.

February 25, 2019

| Support for Amazon Storage Gateway Hardware Appliance | The Amazon Storage Gateway Hardware Appliance includes Storage Gateway software preinstalled on a third-party server. You can manage the appliance from the Amazon Web Services Management Console. The appliance can host file, tape, and Volume Gateways. For more informati on, see <u>Using the Storage</u> Gateway Hardware Appliance . | September 18, 2018 |
|---|--|--------------------|
| Support for Server Message<br>Block (SMB) protocol    | File Gateways added support for the Server Message Block (SMB) protocol to file shares. For more information, see <a href="Creating a File Share">Creating a File Share</a> .  | June 20, 2018      |

#### **Earlier updates**

The following table describes important changes in each release of the *Amazon Storage Gateway User Guide* before May 2018.

| Change   | Description   | Date Changed  |
|--|---|---------------|
| Support for S3<br>One Zone-IA<br>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 <a href="Mailto:Create a file share">Create a file share</a> . | April 4, 2018 |
| New Amazon Web<br>Services Region              | Tape Gateway is now available in the Asia Pacific (Singapore) Region. For detailed information, see <a href="Supported Amazon Regions">Supported Amazon Regions</a> .   | April 3, 2018 |

| Change   | Description   | <b>Date Changed</b>   |
|--|---|-----------------------|
| Support for refresh cache notification, Requester Pays, and canned ACL s for Amazon S3 buckets | 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 Storage Gateway API Reference.  For File Gateways, you can now specify that the requester or reader pays for access charges instead of the bucket owner.  File Gateway can now give full control of written files to the owner of the S3 bucket that maps to the NFS file share. | March 1, 2018         |
| New Amazon Web<br>Services Region  | Storage Gateway is now available in the Europe (Paris) Region. For detailed information, see <a href="Supported Amazon Regions">Supported Amazon Regions</a> .  | December 18,<br>2017  |
| Support for file upload notificat ion and guessing of the MIME type                            | File Gateway can now send notifications when all files written to your NFS file share have been uploaded to Amazon S3. For more information, see <a href="NotifyWhenUploaded">NotifyWhenUploaded</a> in the Storage Gateway API Reference.  File Gateway can now guess the MIME type for uploaded objects based on file extensions. For more information, see <a href="Create a file share">Create a file share</a> .                                     | November 21,<br>2017  |
| Support for<br>VMware ESXi<br>Hypervisor version<br>6.5  | Amazon 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,<br>2017 |

| Change  | Description   | <b>Date Changed</b> |
|---|---|---------------------|
| File Gateway<br>support for<br>Microsoft Hyper-V<br>hypervisor            | You can now deploy a File Gateway on a Microsoft Hyper-V hypervisor. For information, see <a href="Supported hypervisors and host requirements">Supported hypervisors and host requirements</a> .   | June 22, 2017       |
| New Amazon Web<br>Services Region   | Storage Gateway is now available in the Asia Pacific (Mumbai) Region. For detailed information, see <a href="Supported Amazon Regions">Supported Amazon Regions</a> .   | May 02, 2017        |
| Updates to file share settings  Support for cache refresh for file shares | File Gateways now add mount options to the file share settings. You can now set squash and readonly options for your file share. For more information, see <a href="Create a file share">Create a file share</a> .  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 <a href="RefreshCache">RefreshCache</a> in the API Reference.   | March 28, 2017      |
| Support for File<br>Gateways on<br>Amazon EC2                             | Amazon 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 <a href="Create and activate an Amazon S3">Create and activate an Amazon S3</a> File Gateway. For information about how to launch a File Gateway AMI, see <a href="Deploying an Amazon EC2">Deploying an Amazon EC2</a> instance to host an S3 File Gateway.  In addition, File Gateway now supports HTTP proxy configuration. For more information, see <a href="Routing">Routing</a> your gateway deployed on Amazon EC2 through an HTTP proxy. | February 08, 2017   |

| Change                            | Description   | Date Changed         |
|-----------------------------------|---|----------------------|
| New Amazon Web<br>Services Region | Storage Gateway is now available in the Europe (London) Region. For detailed information, see <a href="Supported Amazon Regions">Supported Amazon Regions</a> .   | December 13,<br>2016 |
| New Amazon Web<br>Services Region | Storage Gateway is now available in the Canada (Central) Region. For detailed information, see <a href="Supported Amazon Regions">Supported Amazon Regions</a> .  | December 08,<br>2016 |
| Support for File<br>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,<br>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 <u>DescribeGatewayInformation</u> API action using an Amazon 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.



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 <a href="Managing Gateway Updates Using the Amazon Storage Gateway Console">Managing Gateway Updates Using the Amazon Storage Gateway Console</a>.

| Release Date | Software Version | Release Notes   |
|--------------|------------------|---|
| 2024-04-01   | 1.24.4           | <ul> <li>Addressed missing Network<br/>Time Protocol (NTP)<br/>component</li> </ul>     |
| 2024-03-18   | 1.24.3           | <ul> <li>Fixed an issue related<br/>to case-sensitive lookup<br/>performance</li> </ul> |

| Release Date | Software Version | Release Notes  |
|--------------|------------------|--|
|              |                  | <ul> <li>Fixed an issue that caused processes to crash</li> <li>Updated operating system and software elements to address performance and security concerns</li> </ul> |
| 2024-01-12   | 1.24.2           | <ul> <li>Addressed SMB stability issue</li> </ul>  |
| 2023-12-27   | 1.24.1           | <ul> <li>Addressed SMB stability issue</li> </ul>  |

| Release Date | Software Version | Release Notes   |
|--------------|------------------|---|
| 2023-12-01   | 1.24.0           | <ul> <li>Updated SMB stack</li> <li>Added support for<br/>AES-256 encryption, plus<br/>more secure variants of<br/>AES-128 encryption and<br/>signing when using an<br/>SMB 3.1.1 client which<br/>requests it.</li> <li>SMBv1 (LANMAN/CIFS)<br/>server-side copy and<br/>server-side wildcard<br/>expansion functiona<br/>lity have been removed.<br/>(SMBv2 and SMBv3<br/>are unaffected.) This<br/>may negatively impact<br/>performance of certain<br/>SMBv1 workloads. If<br/>you use SMBv1, you are<br/>encouraged to migrate to<br/>SMBv2 or SMBv3.</li> <li>Updated operating system<br/>and software elements to<br/>address performance and<br/>security concerns.</li> </ul> |
| 2023-10-24   | 1.23.2           | <ul> <li>Fixed an issue related<br/>to Support Channel not<br/>connecting properly for<br/>certain users.</li> </ul>  |
| 2023-08-14   | 1.23.1           | <ul> <li>Updated NTP server to<br/>use sync server for new<br/>gateways.</li> </ul>   |

| Release Date | Software Version | Release Notes  |
|--------------|------------------|--|
| 2023-06-12   | 1.23.0           | <ul> <li>Fixed an access violation issue on large copies.</li> <li>Fixed an NFS issue.</li> <li>Removed Java 8.</li> <li>Increased upload threads for some Amazon accounts.</li> <li>Updated operating system and software elements to address performance and security concerns.</li> </ul> |
| 2023-04-19   | 1.22.1           | <ul> <li>Fixed an issue related to<br/>renaming folders and files.</li> </ul>  |
| 2023-01-18   | 1.22.0           | <ul> <li>Added support for DOS attributes.</li> <li>Increased the number of supported file shares per gateway from 10 to 50.</li> <li>Implemented a clock skew detection mechanism to determine when a gateway and a service are out of sync.</li> <li>Updated SMB stack.</li> </ul>         |
| 2022-07-06   | 1.21.2           | <ul> <li>Updated operating system<br/>and software elements to<br/>address performance and<br/>security concerns.</li> </ul>   |

| Release Date | Software Version | Release Notes  |
|--------------|------------------|--|
| 2022-02-16   | 1.21.1           | <ul> <li>Added new metrics for<br/>rename and deletion in<br/>cache.</li> <li>Fixed other miscellaneous<br/>issues.</li> </ul> |
| 2022-01-18   | 1.21.0           | <ul> <li>Added new CloudWatch<br/>metrics.</li> <li>Added bandwidth throttlin<br/>g for data uploads.</li> </ul>               |
| 2021-12-12   | 1.20.0           | <ul> <li>Addressed Log4j vulnerabi<br/>lity.</li> </ul>  |