How do I secure the files in my Amazon S3 bucket?

By AWS Official

Published: 2017-07-25 · Archived: 2026-04-05 21:16:57 UTC

I want to secure my Amazon S3 bucket with access restrictions, resource monitoring, and data encryption to protect my files and meet security best practices.

Resolution

First, identify whether your Amazon S3 bucket type is general purpose, directory, or table. Then, choose the security measures and monitoring services that align with your bucket type.

Restrict access to your S3 resources

By default, all S3 buckets are private. Only the users who you explicitly grant bucket permissions to can access the bucket.

To restrict access to your S3 buckets or objects, take the following actions:

- Use identity-based policies that specify the users who can access specific buckets and objects. To create and test user policies, use the AWS Policy Generator and IAM policy simulator.
- Use bucket policies that define access to specific buckets and objects. Use a bucket policy to grant access across AWS accounts, grant public or anonymous permissions, and allow or block access based on conditions.
  Note: You can use a **Deny** statement in a bucket policy to restrict access to specific AWS Identity and Access Management (IAM) users even when you granted access to the users in an IAM policy.
- Use Amazon S3 Block Public Access as a centralized way to limit public access. Block Public Access settings override bucket policies and object permissions. Be sure to turn on Block Public Access for all accounts and buckets that you don't want publicly accessible. Amazon S3 turns on Block Public Access by default for all new accounts and buckets. Turn off the feature only when you explicitly require public access to your S3 resources. If you turn off Block Public Access on a bucket, then regularly audit the bucket.
- Set access control lists (ACLs) on your buckets and objects.
  Note: If you must programmatically manage permissions, then use IAM policies or bucket policies instead of ACLs. However, you can use ACLs when your bucket policy exceeds the 20 KB maximum file size. Or, you can use ACLs to grant access for Amazon S3 server access logs or Amazon CloudFront logs.
- Use service control policies (SCPs) to centrally manage and enforce S3 security policies across all accounts in your organization.
- At the network level, restrict access with virtual private cloud (VPC) endpoints, IP address-based restrictions in bucket policies, and AWS PrivateLink for S3. VPC endpoints allow private access to Amazon S3 without internet access.

- Use S3 Access Points to simplify security management for buckets that multiple applications or teams access.
- Implement S3 Object Lock so that users can't delete or overwrite objects within a specified time frame.

If you use ACLs to secure your resources, then implement the following best practices:

- Review ACL permissions that allow Amazon S3 actions on a bucket or an object.
- Restrict who gets **Read** and **Write** access to your buckets.
- Grant **Read** access to the **Everyone** group only when you want everyone to access the bucket or object.
- Don't grant **Write** access to the **Everyone** group. Anyone who has **write** access can add objects to your bucket, and AWS charges you for every uploaded object. Also, anyone with **write** access can delete objects in the bucket.
- Don't grant **Write** access to the **Any authenticated AWS user** group because it includes anyone with an active account. To control access for IAM users on your account, use an IAM policy instead. For more information about how Amazon S3 evaluates IAM policies, see How Amazon S3 authorizes a request.
- For new buckets, Amazon S3 sets S3 Object Ownership to **Bucket owner enforced** by default. This turns off ACLs. To maintain full control over all objects, it's a best practice to turn off ACLs and use bucket policies and IAM policies for access control.

You can also restrict access to specific actions in the following ways:

- To require users to use multi-factor authentication before they can delete an object or turn off bucket versioning, configure MFA delete.
- Set up MFA-protected API access so that uses must authenticate with an AWS MFA device before they call certain Amazon S3 API operations.
- If you temporarily share an S3 object with another user, then create a presigned URL to grant time-limited access to the object.

Monitor your S3 resources

To turn on logging and monitor your S3 resources, take the following actions:

- Activate AWS CloudTrail logging for objects in a bucket. By default, CloudTrail monitors only bucket-level actions. To monitor object-level actions, such as **GetObject**, log data events. For examples of data events, see Examples: Logging data events for Amazon S3 objects.
- Turn on Amazon S3 server access logging. For information about how to review server access logs, see Amazon S3 server access log format.
- Use AWS Config to monitor bucket ACLs and bucket policies for violations that allow public read or write access. For more information, see s3-bucket-public-read-prohibited and s3-bucket-public-write-prohibited.
- Use IAM Access Analyzer to review bucket or IAM policies that grant access to your S3 resources from another account.
- Turn on Amazon Macie to automate the identification of sensitive data that's stored in your buckets, broad access to your buckets, and unencrypted buckets in your account.
- Use CloudTrail with other AWS services to invoke specific processes when you take specific actions on your S3 resources. For example, you can use Amazon EventBridge to log S3-object level operations.

- Use the S3 bucket permissions check from AWS Trusted Advisor to notify you about buckets with open access permissions. For more information, see the AWS Trusted Advisor check reference.

Use encryption to protect your data

If you require encryption during transmission, then use HTTPS protocol to encrypt data in transit to and from Amazon S3. All AWS SDKs and AWS tools use HTTPS by default.

Note: If you use third-party tools to interact with Amazon S3, then contact the third-party company to confirm that their tools also support the HTTPS protocol.

If you require encryption for data at rest, then use the server-side encryption (SSE) options Amazon S3 managed keys (SSE-S3), AWS Key Management Service (AWS KMS) keys (SSE-KMS), or customer-provided keys (SSE-C). SSE provides an additional layer of protection and detailed audit trails through CloudTrail. You can specify the SSE parameters when you write objects to the bucket. You can also turn on default encryption on your bucket with SSE-S3 or SSE-KMS.

Note: Amazon S3 automatically turns on SSE-S3 for all new buckets.

If you require client-side encryption, then see Protecting data by using client-side encryption.

Related information

Identity and access management in Amazon S3

Data protection in Amazon S3

How do I require users from other AWS accounts to use MFA to access my Amazon S3 buckets?

How do I see who accessed my Amazon S3 buckets and objects?
