AMAZON SIMPLE STORAGE SERVICE (S3)

Amazon S3 Security and Data Protection

Tools and Best Practices

Table of Contents

Introduction

Security is a shared responsibility

Security is the top priority at AWS, and Amazon S3 is secure by default. As an AWS customer, you benefit from an advanced global infrastructure that is monitored 24/7 and designed to meet the security requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. AWS is responsible for the “security of the cloud,” whereas customers are responsible for “security in the cloud:"

- **Security of the cloud**
  AWS is responsible for protecting the infrastructure that runs Amazon Simple Storage Service (S3). The effectiveness of our security is regularly tested and verified by third-party auditors as part of the AWS compliance programs.
- **Security in the cloud**
  You are responsible for managing access to your data by using tools to apply the appropriate permissions and access levels. You are also responsible for ensuring that your use of Amazon S3 meets your organization's requirements, including meeting any applicable laws and regulations.

This eBook describes best practices that you should consider as you develop and implement your own security policies. This guidance addresses the foundational aspects of Amazon S3 security and data protection and highlights things you should evaluate in the context of your own environment.

Security and data protection tools and best practices

Access management

- Amazon S3 automatically enables S3 Block Public Access and disables S3 access control lists (ACLs) for all new S3 buckets.
- For existing buckets, turn on S3 Block Public Access to block all public access, and use S3 Object Ownership to disable ACLs.
- Use AWS Identity and Access Management (IAM) roles for applications and AWS services that require Amazon S3 access.
- Amazon S3 automatically encrypts all new objects, applying S3 managed server-side encryption (SSE-S3) as a base level of encryption to all new objects. You can choose to update this default configuration using server-side encryption with AWS Key Management Service keys (SSE-KMS), dual-layer server-side encryption with keys stored in AWS KMS (DSSE-KMS), or server-side encryption with customer-provided keys (SSE-C).
- Use Amazon S3 Virtual Private Cloud (VPC) endpoints and service control policies to limit access and permissions.
- Limit access to specific VPCs or known IP address ranges with bucket policies and access point policies.
- Consider Amazon S3 presigned URLs or Amazon CloudFront signed URLs to provide limited-time access to Amazon S3 for specific applications.
- Within AWS, all data is encrypted in-transit. Use TLS1.2 or higher to enforce encryption in transit for access to Amazon S3.

Monitor and audit security settings

- Use AWS IAM Access Analyzer for S3 to monitor access to your data.
- Use Amazon S3 Inventory to audit and report on the replication and encryption status of your objects for business, compliance, and regulatory needs. In addition, use S3 Inventory to review all of the object ACLs in your buckets before migrating to IAM-based bucket policies and disabling ACLs.
- Audit Amazon S3 API actions using AWS CloudTrail.
- Monitor data access from Amazon S3 with S3 server access logging.
- Use S3 server access logs or AWS CloudTrail to identify requests that rely upon on ACL authorization to succeed before migrating to IAM-based bucket policies and disabling ACLs.
- Use Amazon S3 Storage Lens to identify buckets that aren't following data protection best practices, such as using S3 Replication or S3 Versioning.
- Use Amazon GuardDuty for S3 for threat detection to protect your data.
- Use Amazon Macie to discover and protect your sensitive data at scale.
- Use AWS Trusted Advisor for recommendations to help close security gaps.

Data protection

- Enable object versioning with S3 Versioning.
- Enable Multi-factor Authentication (MFA) Delete and S3 Object Lock when appropriate.
- Consider S3 Replication, including across different AWS Regions and accounts, to protect your data.
- Have a backup strategy in place.

Access management

By default, all Amazon S3 resources—buckets, objects, and related sub-resources (for example, lifecycle configuration and website configuration)—are private: only the resource owner, the AWS account that created the resource, can access the resource. The resource owner can optionally grant access permissions to others by writing an access policy.

Amazon S3 offers access policy options broadly categorized as resource-based policies and user-based policies. Resource-based policies are access policies that you attach to your resources (buckets and access points). User-based policies are access policies attached to users in your account. You can choose to use resource-based policies, user policies, or some combination of these to manage permissions to your Amazon S3 resources.

When granting permissions, you decide who is getting them, the Amazon S3 resources they are applied to, and specific actions you want to allow on those resources.

Amazon S3 default access management settings

S3 Block Public Access is enabled by default for all new S3 buckets

To ensure that public access to all your new Amazon S3 buckets and objects is blocked, S3 Block Public Access is turned on by default for all new buckets. This includes all new buckets regardless of how they are created, including the S3 console, AWS Command Line Interface (CLI), APIs, SDKs, and Amazon CloudFormation. With a few clicks in the S3 console, you can also apply S3 Block Public Access to every bucket in your account or to individual buckets.

AWS recommends that you turn on all Amazon S3 Block Public Access settings. Before applying these settings, ensure that your applications will work correctly without public access. If you require some level of public access to your buckets or objects, you can customize the individual block public access settings to suit your specific storage use cases.

Amazon S3 Block Public Access settings override S3 permissions that allow public access, making it simpler for account administrators to set up centralized controls to prevent variation in security configuration regardless of how an object is added or a bucket is created.

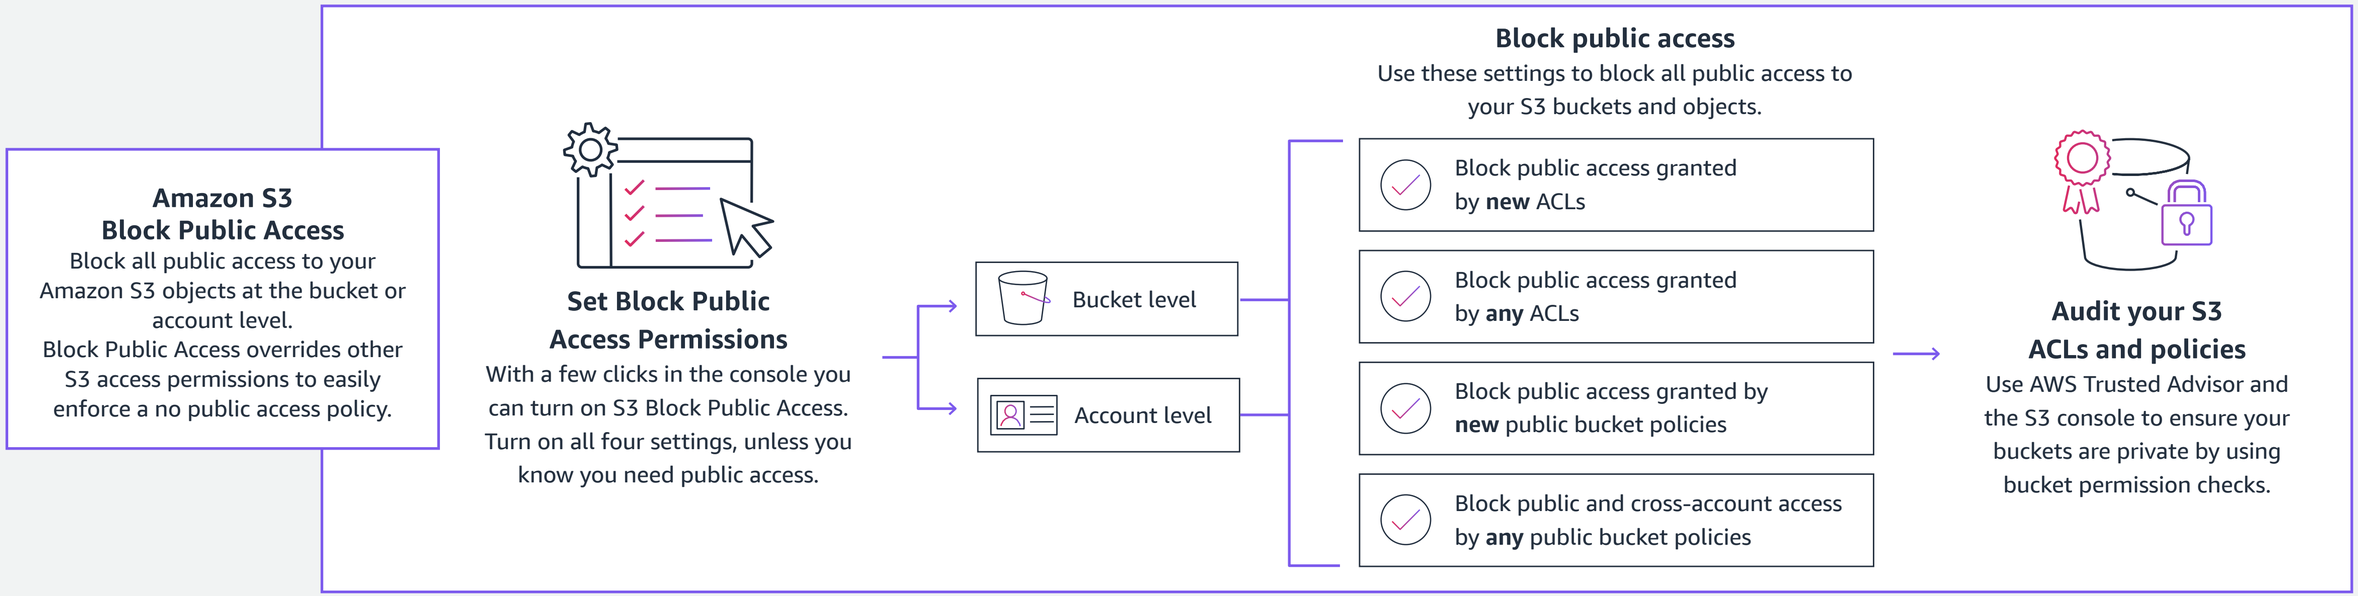

Access control lists (ACLs) are disabled by default for all new S3 buckets

S3 Object Ownership is an Amazon S3 bucket-level setting that you can use to control ownership of objects uploaded to your bucket and to disable or enable access control lists (ACLs). By default, Object Ownership is set to the Bucket owner enforced setting and all ACLs are disabled. When ACLs are disabled, the bucket owner owns all the objects in the bucket and manages access to data exclusively using access management policies.

A majority of modern use cases in Amazon S3 no longer require the use of ACLs, and we recommend that you keep ACLs disabled except in unusual circumstances where you must control access for each object individually. With ACLs disabled, you can use policies to more easily control access to every object in your bucket, regardless of who uploaded the objects in your bucket.

If you would like a better understanding of your ACL usage before migrating to policies and using S3 Object Ownership to disable ACLs, you can use S3 Inventory to review the object ACLs in your buckets and S3 server access logs or AWS CloudTrail to identify requests that rely upon on ACL authorization to succeed.

S3 Object Ownership has three settings that you can use to control ownership of objects uploaded to your bucket and to disable or enable ACLs:

ACLs disabled

- Bucket owner enforced (default): ACLs are disabled, and the bucket owner automatically owns and has full control over every object in the bucket. ACLs no longer affect permissions to data in the S3 bucket. The bucket uses policies to define access control.

ACLs enabled

- Bucket owner preferred: The bucket owner owns and has full control over new objects that other accounts write to the bucket with the bucket-owner-full-control canned ACL.
- Object writer: The AWS account that uploads an object owns the object, has full control over it, and can grant other users access to it through ACLs.

Automatic base encryption with additional options

Amazon S3 automatically encrypts all new objects

Amazon S3 automatically applies S3 managed server-side encryption (SSE-S3) as a base level of encryption to all new objects added to S3, at no additional cost and with no impact on performance. SSE-S3 uses 256-bit Advanced Encryption Standard and has been configured for trillions of objects by customers. This base level of encryption helps customers meet their encryption requirements, with no changes to applications. Alternatively, customers can still choose to update this default configuration using server-side encryption with AWS Key Management Service keys (SSE-KMS), dual-layer server-side encryption with keys stored in AWS KMS (DSSE-KMS), or server-side encryption with customer-provided keys (SSE-C).

To encrypt your existing unencrypted Amazon S3 objects, you can use Amazon S3 Batch Operations. You provide S3 Batch Operations with a list of objects to operate on, and Batch Operations calls the respective API to perform the specified operation. You can use the Batch Operations Copy operation to copy existing unencrypted objects and write them back to the same bucket as encrypted objects. A single Batch Operations job can perform the specified operation on billions of objects.

You can also encrypt existing objects by using the CopyObject API operation or the copy-object AWS CLI command.

Securing data using server-side encryption

Server-side encryption is the encryption of data at its destination by the application or service that receives it. Amazon S3 encrypts your data at the object level as it writes it to disks in its data centers and decrypts it for you when you access it. As long as you authenticate your request and you have access permissions, there is no difference in the way you access encrypted or unencrypted objects.

1. Server-side encryption with Amazon S3 managed keys (SSE-S3)

When you use Server-Side Encryption with Amazon S3-Managed Keys (SSE-S3), each object is encrypted with a unique key. As an additional safeguard, it encrypts the key itself with a key that it rotates regularly. Amazon S3 server-side encryption uses 256-bit Advanced Encryption Standard Galois/Counter Mode (AES-GCM) to encrypt all uploaded objects. For more information, visit the S3 User Guide.

2. Server-side encryption with AWS Key Management Service keys (SSE-KMS)

SSE-KMS lets AWS Key Management Service (AWS KMS) manage your encryption keys. Using AWS KMS to manage your keys provides several additional benefits. With AWS KMS, there are separate permissions for the use of the KMS key, providing an additional layer of control and protection against unauthorized access to your objects stored in Amazon S3. AWS KMS provides an audit trail so you can see who used your key to access which object and when, as well as view failed attempts to access data from users without permission to decrypt the data. Additionally, AWS KMS provides additional security controls to support customer efforts to comply with PCI-DSS, HIPAA/HITECH, and FedRAMP industry requirements. Amazon S3 Bucket Keys can reduce the request costs of SSE-KMS by up to 99% by decreasing the request traffic from S3 to KMS during cryptographic operations. For more information on SSE-KMS, visit the S3 User Guide.

3. Dual-layer server-side encryption with keys stored in AWS Key Management Service (DSSE-KMS)

Dual-layer server-side encryption with keys stored in AWS Key Management Service (DSSE-KMS) is designed to meet National Security Agency CNSSP 15 for FIPS compliance and Data-at-Rest Capability Package (DAR CP) Version 5.0 guidance for two layers of CNSA encryption. Amazon S3 is the only cloud object storage service where customers can apply two independent layers of encryption at the object level and control the data keys used for both layers. Each layer of encryption uses a different implementation of the 256-bit Advanced Encryption Standard with Galois Counter Mode (AES-GCM) algorithm.

4. Server-side encryption with customer-provided keys (SSE-C)

With Server-Side Encryption with Customer-Provided Keys (SSE-C), you manage the encryption keys and Amazon S3 manages the encryption, as it writes to disks, and decryption, when you access your objects. For more information, visit the S3 User Guide.

Enforce encryption of data in transit

To protect data while in transit (as it travels to and from Amazon S3), you can encrypt data in transit using Secure Socket Layer/Transport Layer Security (SSL/TLS) or client-side encryption. You can use condition keys in bucket policies to enforce HTTPS over TLS to help prevent potential unauthorized users from eavesdropping on or manipulating network traffic using person-in-the-middle or similar methods. If you use the AWS Encryption SDK, this is done for you.

“Least privilege” access model is the cornerstone of Amazon S3 security best practices

Following a “least privilege” access control model means only granting users permission to access resources that are absolutely necessary to performing their respective job duties. You should grant only the permissions that are required to perform a task.

It is a best practice to start with no privileges (no permissions) and incrementally add them over time to specific project teams or users that need access to those Amazon S3 resources. It’s an easier and lower-risk method to audit access to your resources, as opposed to starting with an open base of users and denying permissions. Therefore, you should grant only the permissions that are required to perform a task.

Implementing least privilege access is fundamental in reducing security risk and the impact that could result from errors or malicious intent.

How we enable you to implement a “least privilege” access control model:

- AWS Identity and Access Management (IAM) directly enables fine-grained S3 access controls, with no permissions by default.
- Amazon S3 bucket policies restrict bucket and object access by user, network, or application across accounts, and are private by default.
- Amazon S3 Access Points grant different users a separate set of permissions, and can firewall your data by restricting access to a VPC.
- Amazon S3 object tags are metadata you can reference in AWS IAM and S3 bucket policies to control permissions to specific users (e.g., Finance, HR).
- Amazon S3 Block Public Access is enabled by default for all new buckets.
- Access control lists (ACLs) are disabled by default for all new buckets.

“Least privilege” access model: Manage access to shared datasets with Amazon S3 Access Points

Use Amazon S3 Access Points to simplify managing access for shared S3 data sets

Amazon S3 Access Points simplify managing data access at scale for shared datasets in Amazon S3. Each access point has distinct permissions and network controls that Amazon S3 applies for any request that is made through that access point. Each access point enforces a customized access point policy that works in conjunction with the bucket policy that is attached to the underlying bucket.

Configuring IAM policies for using access points

Amazon S3 Access Points support AWS IAM resource policies that allow you to control the use of the access point by resource, user, or other conditions. For an application or user to be able to access objects through an access point, both the access point and the underlying bucket must permit the request.

Managing public access to access points

Amazon S3 access points support independent block public access settings for each access point. When you create an access point, you can specify block public access settings that apply to that access point.

VPC-specific access points

You can configure any access point to accept requests only from a VPC to restrict Amazon S3 data access to a private network.

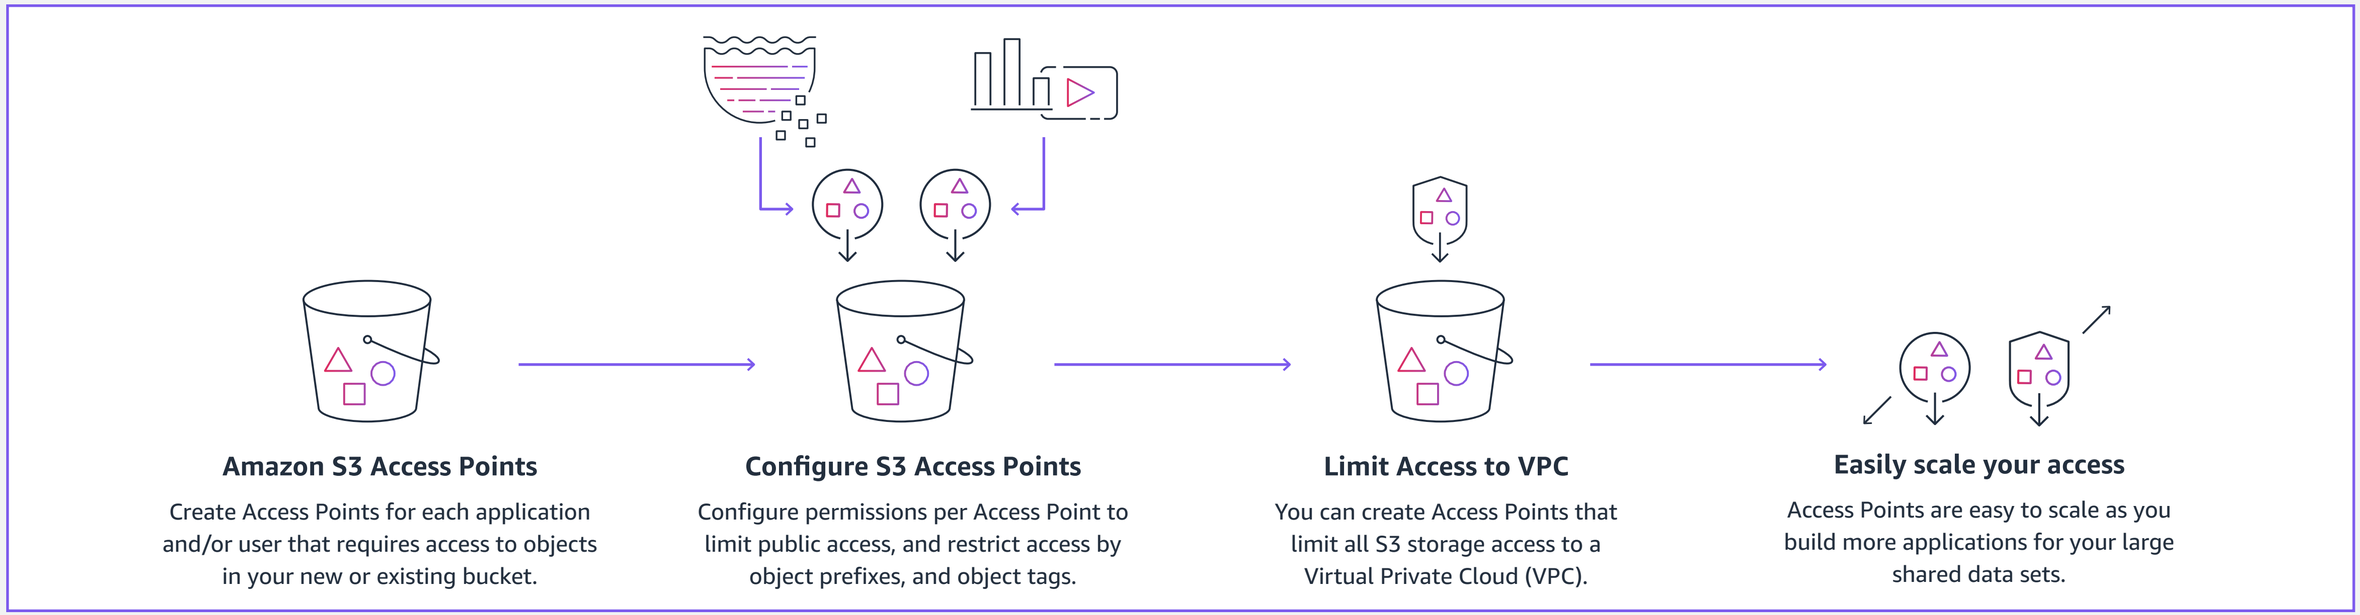

Use IAM roles for applications and AWS services that require S3 access

You can create and configure IAM user or role policies for controlling access to Amazon S3. User or role policies use JSON-based access policy language.

For applications on Amazon Elastic Compute Cloud (Amazon EC2) or other AWS services to access Amazon S3 resources, they must include valid AWS credentials in their AWS API requests. You should not store AWS credentials directly in the application or Amazon EC2 instance.

You should use an IAM role to manage temporary credentials for applications or services that need to access Amazon S3. When you use a role, you don’t have to distribute long-term credentials (such as a user name and password or access keys) to an Amazon EC2 instance or AWS service such as AWS Lambda. The role supplies temporary permissions that applications can use when they make calls to other AWS resources.

AWS Organizations service control policies (SCPs)

Service control policies (SCPs) are a type of organization policy that you can use to manage permissions in your organization. An SCP defines a guardrail, or sets limits, on the actions that the account's administrator can delegate to the IAM users and roles in the affected accounts. The administrator must still attach identity-based or user-based policies to IAM users or roles, or to the resources in your accounts to actually grant permissions. The effective permissions are the logical intersection between what is allowed by the SCP and what is allowed by the IAM and resource-based policies.

Use S3 bucket policies to grant access permissions to your S3 bucket and the objects in it

A bucket policy is a resource-based policy that you can use to grant access permissions to your Amazon S3 bucket and the objects in it. The permissions attached to the bucket apply to all of the objects in the bucket that are owned by the bucket owner. These permissions do not apply to objects that are owned by other AWS accounts.

AWS PrivateLink for Amazon S3

Access Amazon S3 directly as a private endpoint within your secure, virtual network with AWS PrivateLink for S3. Simplify your network architecture by connecting to S3 from on premises or in the cloud using private IP addresses from your VPC. You no longer need to use public IPs, configure firewall rules, or configure an internet gateway to access S3 from on premises.

Audit and monitor Amazon S3 security

Being vigilant about the integrity of your security posture is just as critical as the setup of your access controls and data protection capabilities. AWS has a number of services to help you audit and monitor your security policies.

Identify and audit all of your Amazon S3 buckets

Identification of your IT assets is a crucial aspect of governance and security. You need to have visibility of all your Amazon S3 resources to assess their security posture and take action on potential Issues.

Tag Editor

Use Tag Editor to label security-sensitive or audit-sensitive resources, then use those tags when you need to search for these resources.

Amazon S3 Storage Lens

Amazon S3 Storage Lens is a cloud-storage analytics feature that you can use to gain organization-wide visibility into object-storage usage and activity. You can use S3 Storage Lens metrics to identify cost-optimization opportunities, implement data-protection and security best practices, and improve the performance of application workloads. For example, identify buckets that aren't following data protection best practices, such as using S3 Replication or S3 Versioning.

Amazon S3 Inventory

Use Amazon S3 Inventory to audit and report on the replication and encryption status of your objects for business, compliance, and regulatory needs. In addition, use S3 Inventory to review all of the object ACLs in your buckets before migrating to IAM-based bucket policies and disabling ACLs.

Amazon Macie

With Amazon Macie, you can discover, classify, and protect sensitive stored data. Amazon Macie automatically provides you with a full inventory of your S3 buckets by scanning buckets to identify and categorize the data. It also automatically and continually evaluates bucket-level preventative controls while providing actionable security findings.

Implement monitoring of your S3 environment and bucket policies

Monitoring is an important part of maintaining the security, availability, and performance of Amazon S3 and your AWS solutions.

Use AWS Identity and Access Management (IAM) Access Analyzer for S3

AWS IAM Access Analyzer for S3 lists S3 buckets that are configured to allow access to anyone on the internet or other AWS accounts, including AWS accounts outside of your organization. For each public or shared bucket, you receive findings into the source and level of public or shared access. For example, Access Analyzer for S3 might show that a bucket has read or write access provided through a bucket access control list (ACL), a bucket policy, or an access point policy. Armed with this knowledge, you can take immediate and precise corrective action to restore your bucket access to what you intended.

When reviewing an at-risk bucket in Access Analyzer for S3, you can block all public access to the bucket with a single click using S3 Block Public Access. We recommend that you block all access to your buckets unless you require public access to support a specific use case.

You can also drill down into bucket-level permission settings to configure granular levels of access. For specific and verified use cases that require public access, such as static website hosting or cross-account sharing, you can acknowledge and record your intent for the bucket to remain public or shared by archiving the findings for the bucket. You can revisit and modify these bucket configurations at any time. You can also download your findings as a CSV report for auditing purposes.

How Access Analyzer for S3 works

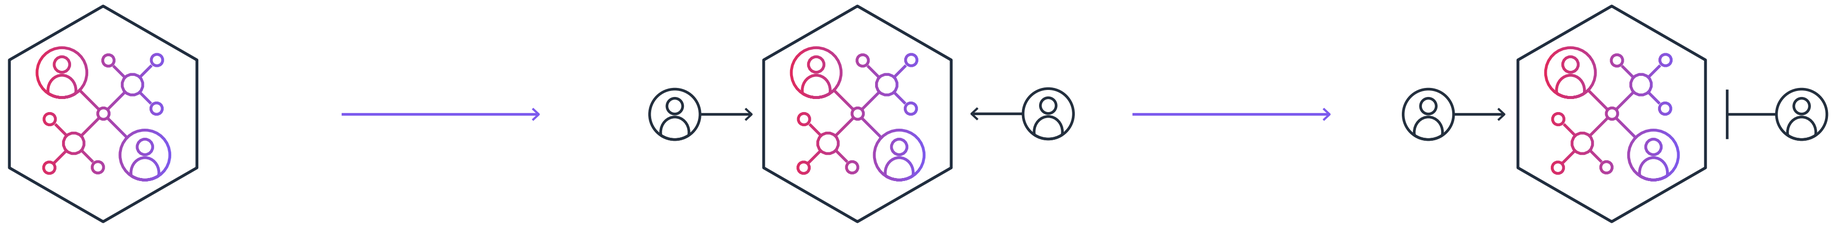

STEP 1 - Create an analyzer

To begin generating AWS IAM Access Analyzer findings, create an analyzer for your account. The analyzer will continuously scan supported resource policies within the account.

STEP 2 - Review your findings

AWS IAM Access Analyzer generates detailed findings for resources that are accessible from outside the AWS account.

STEP 3 - Take action

If access to the resource is not intended, modify the resource policy to further restrict access. If the access to the resource is intended, then you can archive the finding.

Tighten Amazon S3 permissions for your IAM users and roles using access history of Amazon S3 actions

You can use action and access history for users or roles, in combination with AWS IAM Access Analyzer findings, to improve the security posture of your Amazon S3 permissions. To help you identify unused S3 permissions, AWS IAM provides last accessed information for S3 management actions and reports the last time a user or role used an S3 action. This granular access information helps you analyze access, identify unused S3 actions, and remove them confidently.

Use logs to monitor requests and actions in Amazon S3

AWS CloudTrail logs provide a record of actions taken by a user, role, or an AWS service in Amazon S3, including detailed API tracking for S3 bucket-level and object-level operations. S3 server access logs are delivered to a bucket that you choose, and contain details about each request, such as request type, the resource specified in the request, and the time and date the request was processed.

Logs are useful for many applications. For example, access log information can be useful in security and access audits (e.g., you can use S3 server access logs or AWS CloudTrail to identify requests that rely upon on ACL authorization to succeed before migrating to IAM-based bucket policies and disabling ACLs). It can also help you learn about your customer base and understand your S3 bill. By default, bucket logging is not enabled; you should enable logging if you want to perform security audits or learn more about users and usage patterns. For more information about how the different logs work, and their properties, performance and costs, visit the S3 User Guide.

You can also use CloudTrail logs with Amazon CloudWatch to receive alerts or set alarms in response to specific API activity, improving visibility and monitoring.

Use Amazon Macie to gain visibility of your data security posture

Amazon Macie automates the discovery of sensitive data at scale and gives you constant visibility of the data security and data privacy of your data stored in Amazon S3. Macie automatically and continually evaluates all of your S3 buckets and alerts you to any unencrypted buckets, publicly accessible buckets, or buckets shared with AWS accounts outside those you have defined in the AWS Organizations. Macie provides native multi- account support so you can view your data security posture across your entire Amazon S3 environment from a single Macie administrator account.

Macie applies machine learning and pattern matching techniques to the buckets you select to identify and alert you to sensitive data, such as personally identifiable information (PII). This can help you meet regulations, such as the Health Insurance Portability and Accountability Act (HIPAA) and General Data Privacy Regulation (GDPR).

Use Amazon GuardDuty for S3 for threat detection to protect your data

Amazon GuardDuty monitors for suspicious data access and anomaly detection to help you better protect your data residing in Amazon S3. GuardDuty continuously monitors data access behavior and identifies suspicious activity, such as data access from an unusual geo-location, API calls from a known malicious IP address, or unusual API calls consistent with malicious data discovery attempts. For your reference, here's the full list of GuardDuty S3 threat detections.

When threats are detected, Amazon GuardDuty produces detailed security findings to the console and to Amazon EventBridge, making alerts actionable and easy to integrate into existing event management and workflow systems, or trigger automated remediation actions using AWS Lambda. With support for AWS Organizations you can enable Amazon S3 Protection across your entire organization with a single click.

Use AWS Trusted Advisor to inspect your AWS environment

AWS Trusted Advisor inspects your AWS environment and then makes recommendations when opportunities exist to help close security gaps.

AWS Trusted Advisor has the following Amazon S3-related checks:

- Logging configuration of Amazon S3 buckets
- Security checks for Amazon S3 buckets that have open access permissions
- Fault tolerance checks for Amazon S3 buckets that don't have S3 Versioning enabled, or have S3 Versioning suspended

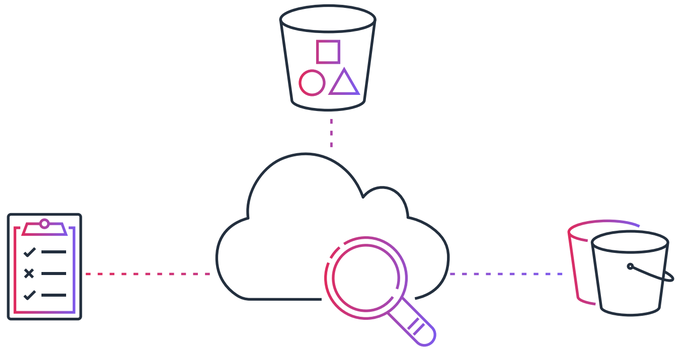

Monitor and manage S3 using Amazon CloudWatch

Amazon CloudWatch metrics for Amazon S3 can help you understand and improve the performance of applications that use Amazon S3.

There are several ways that you can use Amazon CloudWatch with Amazon S3:

- **Daily storage metrics for buckets:** CloudWatch collects and processes storage data from S3 into readable metrics provided daily to all customers at no additional cost.
- **Request metrics:** Minute-level metrics reported for all object operations, either at the bucket-level or for a shared prefix, object tag, or S3 Access Point.
- **Replication metrics:** Monitor the total number of S3 API operations that are pending replication, the total size of objects pending replication, and the maximum replication time to the destination Region.
- **Amazon S3 Storage Lens metrics:** You can publish S3 Storage Lens usage and activity metrics to CloudWatch to create a unified view of your operational health in CloudWatch dashboards.

Protecting data at scale with Amazon S3

A vital function of storage is data protection—primarily protection against corruption, loss, and accidental or malicious overwrites, modifications, or deletions. Amazon S3 has several features, capabilities, and integrations with other AWS services to provide the highest levels of data protection when it is used as the core storage for your applications.

Data protection rests on the durability of the storage platform used. A system that is durable is able to perform its responsibilities over time, even when unexpected events may occur. Amazon S3 is storage infrastructure suitable for mission-critical and primary storage. It is designed to provide 99.999999999% (11 9s) data durability within a single AWS Region.

To help achieve this data durability, Amazon S3 stores your data redundantly across a minimum of three Availability Zones in an AWS Region and on multiple devices in each Availability Zone. This enables Amazon S3 to sustain your data in the event of an entire Availability Zone loss.

As with any environment, the best practice is to have a backup and to put in place safeguards against malicious or accidental deletion. For Amazon S3 data, you can use S3 Versioning, S3 Object Lock, S3 Replication, and functioning, regularly tested backups to fulfill data protection best practices.

Amazon S3 is designed to provide 99.999999999% (11 9s) data durability within an AWS Region.

Use Amazon S3 Versioning to preserve and restore your objects

One key element of data protection is to protect data assets against unintentional and malicious deletion and corruption, whether through users accidentally deleting data assets, applications inadvertently deleting or corrupting data, or outside parties trying to tamper with data.

You can help protect your data be enabling Amazon S3 Versioning, a means of keeping multiple variants of an object in the same bucket. You can use S3 Versioning to preserve, retrieve, and restore every version of every object that is stored in your S3 bucket. With S3 Versioning, you can more easily recover from both unintended user actions, like accidental deletions or overwrites, and application failures. You can use S3 Versioning for data protection and retention scenarios such as recovering objects that have been accidentally deleted or overwritten and archiving previous versions of objects to the Amazon S3 Glacier storage classes for long-term low-cost storage.

When working with Amazon S3 Versioning in S3 buckets, you can optionally add another layer of security by configuring a bucket to enable MFA (multi-factor authentication) delete. By default, all requests to your S3 bucket require your AWS account credentials. If you enable S3 Versioning with MFA delete on your S3 bucket, two forms of authentication are required to permanently delete a version of an object: your AWS account credentials and a valid six-digit, one-time use code and serial number from an authentication device in your physical possession. MFA delete can help prevent accidental bucket deletions. If MFA delete is not enabled, any user with the password of a sufficiently privileged root or IAM user could permanently delete an S3 object.

MFA Delete requires additional authentication for either of the following operations:

- Changing the versioning state of your bucket
- Permanently deleting an object version

To learn more about enabling S3 Versioning with MFA delete, including how to purchase and activate an authentication device, refer to the S3 User Guide.

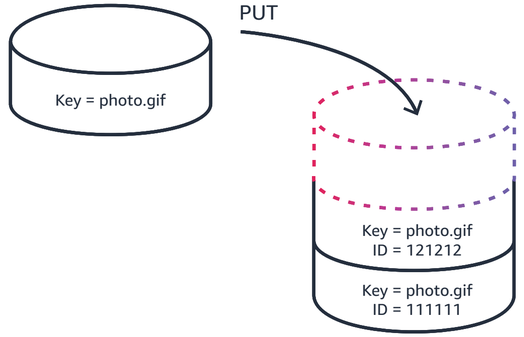

Create a new version with every upload

Previous versions are retained, not overwritten

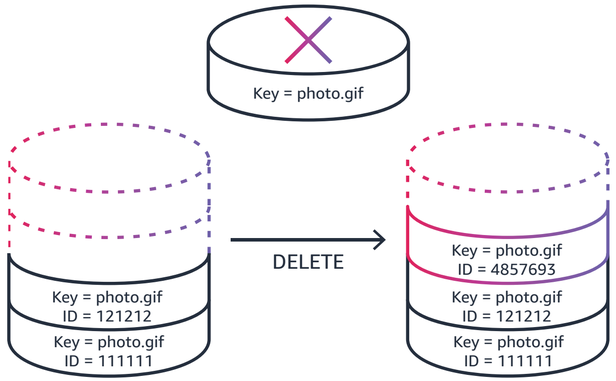

Protect from unintended user delete

Delete requests without version ID removes access to object but keeps the data

Use Amazon S3 Object Lock to enforce retention policies

Amazon S3 Object Lock is the industry standard for object storage immutability for ransomware protection. S3 Object Lock blocks permanent object deletion during a customer-defined retention period so that you can enforce retention policies as an added layer of data protection or for regulatory compliance. With S3 Object Lock, S3 Versioning is automatically enabled, and these features work together to prevent locked object versions from being permanently deleted (accidental or intentional) or overwritten using a write-once-read-many (WORM) model.

Managing object retention with S3 Object Lock

S3 Object Lock provides two ways to manage object retention: *retention periods* and *legal holds*. With S3 Object Lock enabled on a bucket, an object version can have both a retention period and a legal hold, one but not the other, or neither.

- **Retention period:** Specifies a fixed period of time during which an object remains locked. During this period, your object is protected by the WORM configuration and can't be overwritten or deleted. When you place a retention period on an object version, Amazon S3 stores a timestamp in the object version's metadata to show when the retention period expires. After the retention period expires, the object version can be overwritten or deleted unless you also placed a legal hold on the object version. For more information, visit the S3 User Guide.

- **Legal hold:** Provides the same protection as a retention period, but it has no expiration date. Instead, a legal hold remains in place until you explicitly remove it. Legal holds are independent from retention periods. For more information, visit the S3 User Guide.

Retention periods and *retention modes* are always configured in tandem, unlike legal holds, which are configured independently. S3 Object Lock provides two retention modes that apply different levels of protection to your objects. You can apply either retention mode to any object version that is protected by Object Lock.

- **Governance mode:** In governance mode, users can't overwrite or delete an object version or alter its lock settings unless they have special permissions. With governance mode, you protect objects against being deleted by most users, but you can still grant some users permission to alter the retention settings or delete the object if necessary. You can also use governance mode to test retention-period settings before creating a compliance-mode retention period.

- **Compliance mode:** In compliance mode, a protected object version can't be overwritten or deleted by any user, including the root user in your AWS account. When an object is locked in compliance mode, you cannot change the retention mode, and you cannot shorten the retention period. Compliance mode helps ensure that an object version can't be overwritten or deleted for the duration of the retention period. S3 Object Lock has been assessed for SEC Rule 17a-4(f), FINRA Rule 4511, and CFTC Regulation 1.31 by Cohasset Associates.

Replicate resources to meet requirements

Amazon S3 Replication to protect your data and meet compliance requirements

Amazon S3 is designed to provide 99.999999999% (11 9s) data durability within an AWS Region, but many enterprise organizations may have compliance and risk models that require replication of data assets to a second geographically distant location. Amazon S3 Replication enables automatic, asynchronous copying of objects across S3 buckets. With S3 Replication, you can replicate new objects written to a bucket to one or more destination buckets in the same AWS Region or different AWS Regions within the same account or across accounts. You can also replicate existing bucket contents (S3 Batch Replication), including existing objects, objects that previously failed to replicate, and objects replicated from another source.

Recommendations for S3 Replication

- **Replicate objects while retaining metadata**: If you need to ensure your replica copy is identical to the source copy, you can use replication to make copies that retain all metadata.
- **Replicate objects to more cost-effective storage classes**: For long-term backup and archival, you can use Amazon S3 Replication to put objects into the Amazon S3 Glacier storage classes to save costs.
- **Maintain object copies under different ownership**: Regardless of who owns the source object, you can tell Amazon S3 to change replica ownership to the AWS account that owns the destination bucket to restrict access to object replicas.

How S3 Replication works

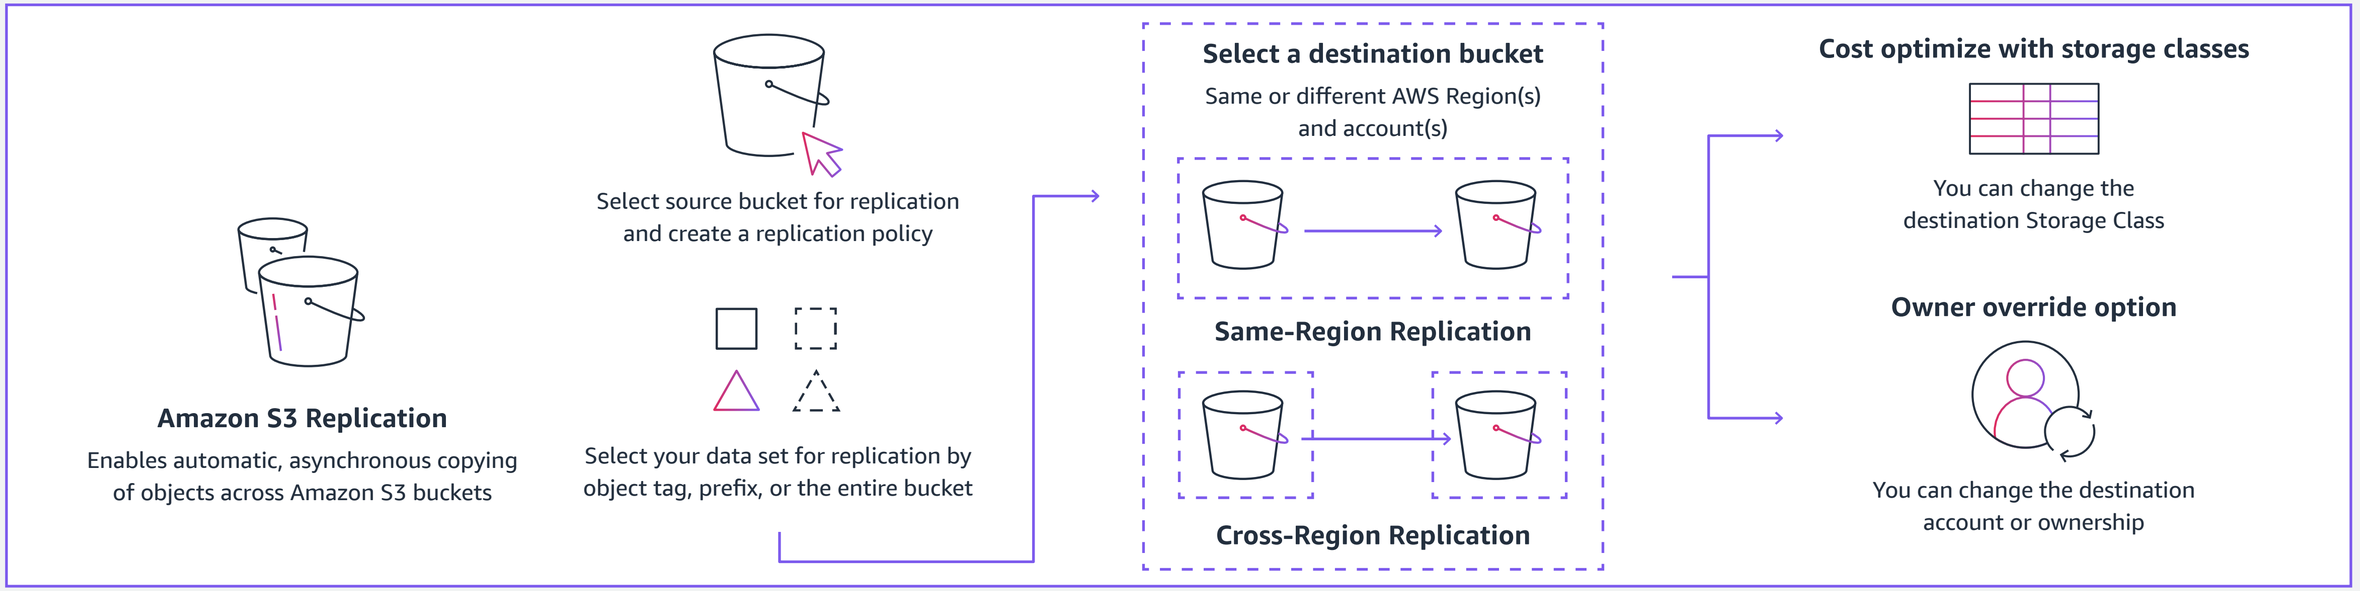

Amazon S3 Replication Time Control

Amazon S3 Replication Time Control (RTC) helps you meet compliance or business requirements for data replication and provides visibility into S3 Replication activity. S3 RTC is backed by a SLA on the replication of 99.9% of objects within 15 minutes during any billing month.

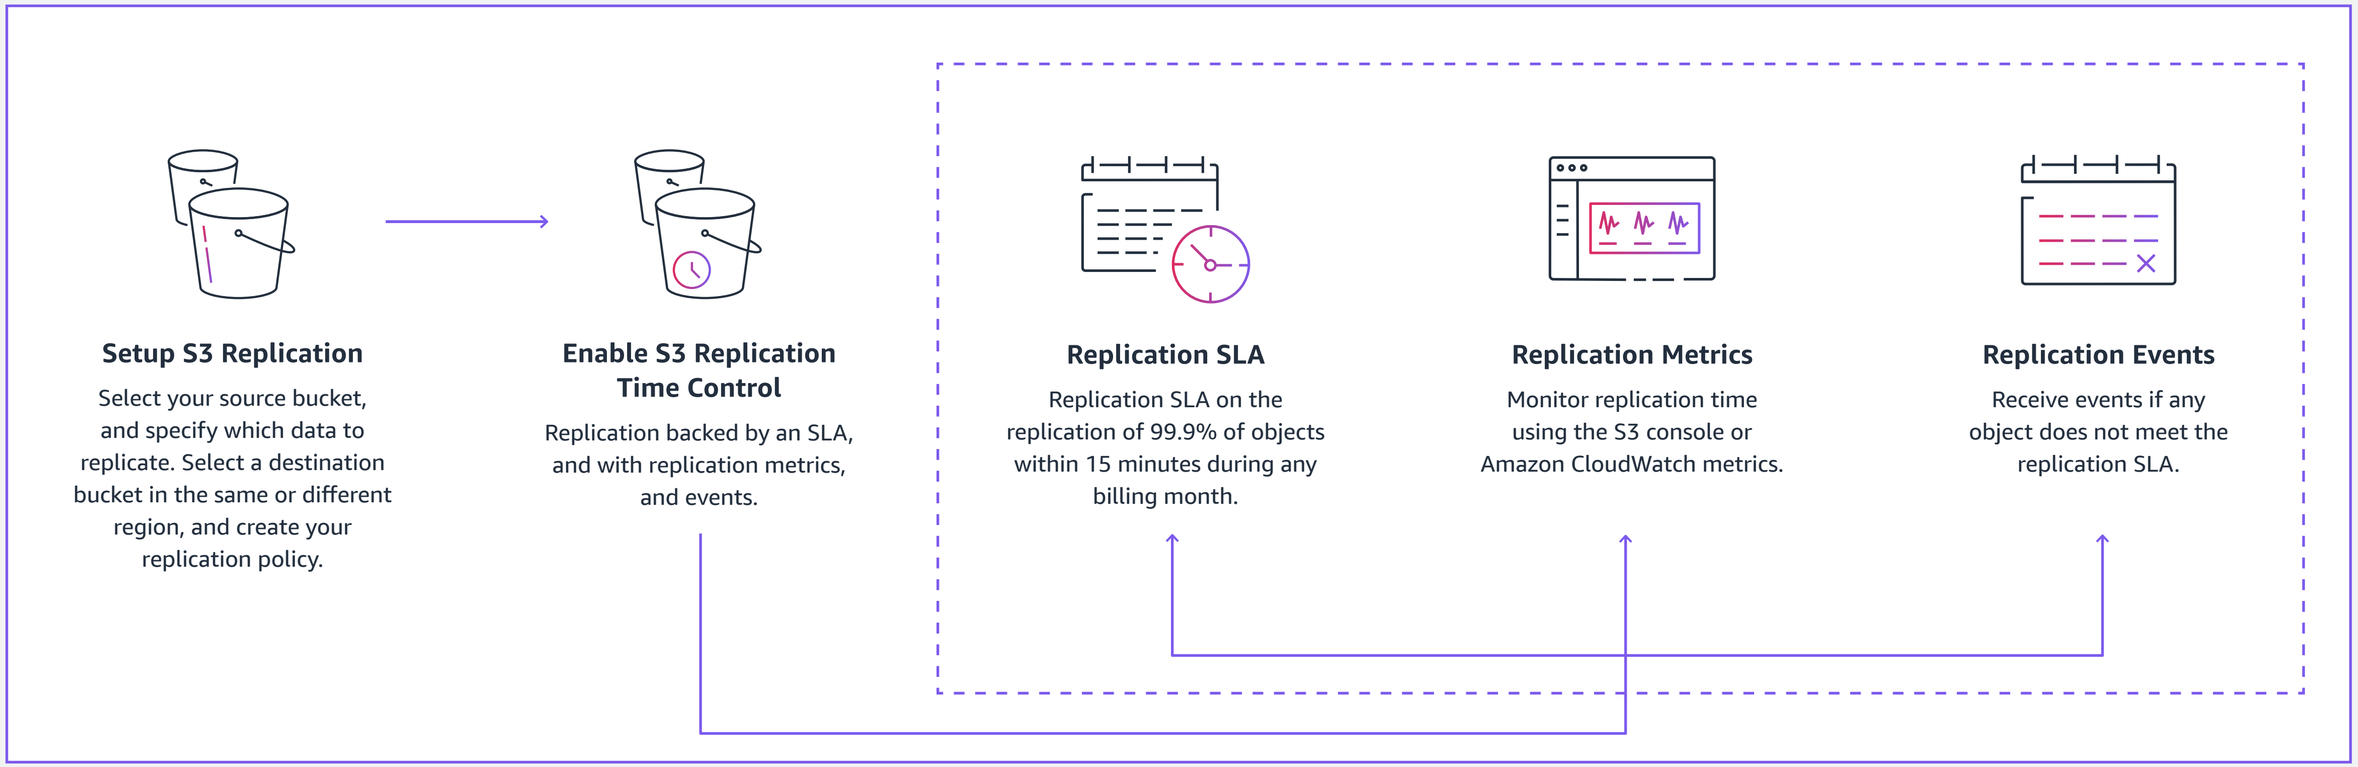

Have a data backup strategy in place

A data backup is a copy of your system, configuration, or application data that's stored separately from the original. Sometimes organizations may experience unexpected events like natural disasters, human errors, security events, or system failures. Data backup is a critical data protection function to decrease the risk of full or partial data loss in the case of unexpected events. It offers organizations the ability to restore systems and applications to a previously desired state.

Amazon Web Services (AWS) and AWS Storage Competency Partners in the AWS Partner Network (APN) can help you build secure, efficient, and cost-effective backup and restore options using tools you know and trust, with the scale and reliability of AWS. Working with these partners can help your organization (particularly if you're managing legacy infrastructure and applications) deploy capabilities that send data to the AWS Cloud for durable storage and protection.

Another customer resource is the AWS Marketplace, an online catalogue offering over 3,500 software listings from over 1,100 independent software vendors, where you can explore, purchase, and deploy off-the-shelf cloud storage solutions. All listed solutions support a variety of AWS storage services, and are ready for deployment and immediate use.

AWS Backup supports Amazon S3, providing, a fully managed, policy-based service that you can use to centrally define backup policies to protect your data in S3. After you define your backup policies and assign S3 resources to the policies, AWS Backup automates the creation of S3 backups and securely stores the backups in an encrypted backup vault that you designate in your backup plan.

Prerequisites

You must activate S3 Versioning on your bucket before AWS Backup can back it up.

When using AWS Backup for Amazon S3, you can perform the following actions:

- Create continuous backups and periodic backups. Continuous backups are useful for point-in-time restore, and periodic backups are useful to meet your long-term data-retention needs.
- Automate backup scheduling and retention by centrally configuring backup policies.
- Restore backups of Amazon S3 data to a point in time that you specify.
- Along with AWS Backup, you can use S3 Versioning and S3 Replication to help recover from accidental deletions and perform your own self-recovery operations.

NOTE

We recommend that you set an Amazon S3 Lifecycle expiration rule for buckets with object versioning enabled that are being backed up. If you do not set an expiration rule, your S3 storage costs might increase because AWS Backup retains all versions of your S3 data.

Conclusion

Takeaways for Amazon S3 Security and Data Protection

Foundational tenets of security and data protection

- Operate under a “Least Privilege” access model and continually review access
- Be vigilant monitoring and auditing your data and security settings
- Protect your data for recovery and to help meet regulatory and internal compliance

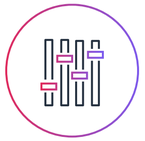

Access controls

- Block all public access using S3 Block Public Access
- Disable all access control lists (ACLs)
- Implement the encryption option that best suits your needs, and enforce encryption of data in transit
- Incrementally grant permissions
- Use AWS IAM policies and S3 bucket policies to manage access to your S3 resources
- Use AWS IAM roles to delegate access to users, applications, or services that don't normally have access to your AWS resources
- Use S3 Access Points to scope application permissions
- Enable VPC endpoints with bucket policies limiting access

Audit and monitor

- Continuously monitor buckets to verify that they meet your security standards
- Scan Amazon S3 to identify public buckets
- Track and limit who is trying to access your Amazon S3 buckets
- Monitor your security settings using AWS tools

Data protection

- Evaluate where S3 Versioning, MFA Delete, S3 Object Lock, and S3 Replication can help you protect your data
- Have a data backup strategy in place

Following these best practices can help you meet stringent data security, compliance, privacy, and protection requirements. Amazon S3 includes a broad range of certifications, including PCI-DSS, HIPAA/HITECH, FedRAMP, SEC Rule 17-a-4, FISMA, EU Data Protection Directive, and many other global agency certifications. These levels of compliance and protection allow organizations to build on Amazon S3 more securely and with less risk than in on-premises data centers.

Additional resources

Amazon S3 security resources

https://aws.amazon.com/s3/security/

Infographic: Configure, automate, and enforce granular access controls for Amazon S3

https://aws.amazon.com/s3/security-infographic/

Documentation: Amazon S3 security

https://docs.aws.amazon.com/AmazonS3/latest/dev/security.html

AWS Storage Blog

https://aws.amazon.com/blogs/storage/