aws storage

EBOOK

Security Best Practices and Guidelines for Amazon S3

Contents

Introduction

Security of Amazon Simple Storage Service (Amazon S3) is a shared responsibility

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from a data center and network architecture that are built to meet the 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 (Amazon 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**

  Your responsibility is for managing access to your data by using tools to apply the appropriate permissions and access levels. You are also responsible for your organization's requirements, and applicable laws and regulations.

This eBook addresses a few foundational Amazon S3 security best practices to consider as you develop and implement your own security policies. The following best practices are general guidelines and don't represent a complete security solution. Because these best practices might not be appropriate or sufficient

Amazon S3 security best practices

for your environment, treat them as helpful considerations rather than prescriptions.

Access control

- Implement a "Least Privilege" access model to limit access to S3 resources by using a combination of Identity and Access Management (IAM) policies, bucket policies, and S3 Access Points
- Ensure that your S3 buckets are not publicly accessible
- Limit access to specific Virtual Private Clouds (VPCs) or known IP address ranges with bucket policies, and access point policies
- Use IAM roles for applications and AWS services that require Amazon S3 access
- Consider Amazon S3 presigned URLs or Amazon CloudFront signed URLs to provide limited-time access to Amazon S3 for specific applications
- Use Amazon S3 VPC Endpoints and Service Control Policies
- Use Access Analyzer for S3 to monitor and control access to your data

Data protection

- Encrypt all Amazon S3 data at rest using Server-side Encryption (SSE) or client-side encryption
- Enforce encryption-in-transit for access to Amazon S3
- Enable object versioning
- Enable Multi-factor Authentication (MFA) Delete and S3 Object Lock when appropriate
- Consider S3 Replication to different AWS accounts to protect your data and remain compliant
- Use tools including Amazon Macie, Amazon GuardDuty for S3, and Amazon S3 Inventory to protect your sensitive data

Monitor and audit security settings

- Audit Amazon S3 API actions using AWS CloudTrail
- Monitor data access from Amazon S3 with access logging

Access Control

By default, all Amazon S3 resources—buckets, objects, and related subresources (for example, lifecycle configuration and website configuration)—are private: only the resource owner, an AWS account that created it, 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 policies. Access policies you attach to your resources (buckets and objects) are referred to as resource-based policies.

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

“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:

- 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 is metadata you can reference in AWS IAM and S3 bucket policies to control permissions to specific users (e.g., Finance, HR).
- Amazon S3 access control lists (ACLs) allow permission, but do not explicitly deny them. Most access control requirements are met using IAM roles, bucket policies and Access Points. There are some limited use cases that require the use of ACLs, as documented here: https://aws.amazon.com/blogs/security/iam-policies-and-bucket-policies-and-acls-oh-my-controlling-access-to-s3-resources/

Use Amazon S3 Access Points to easily manage 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. You can also configure custom block public access settings for each access point.

How S3 Access Points Work

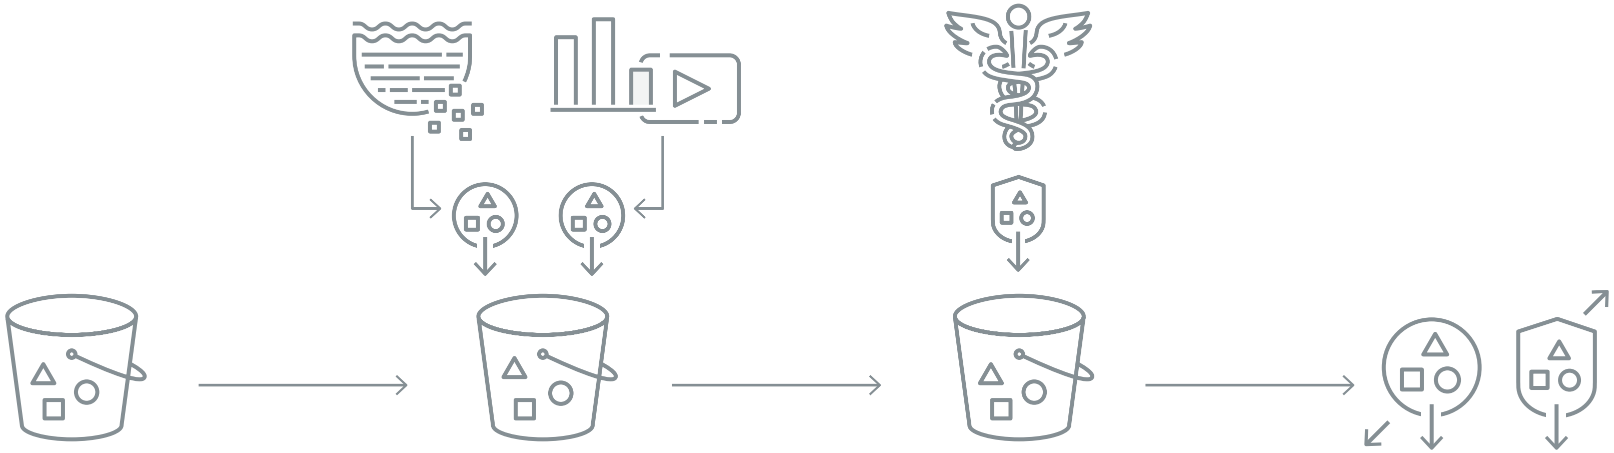

Amazon S3 Access Points

Create Access Points for each application and/or user that requires access to objects in your new or existing bucket

Configure S3 Access Points

Configure permissions per Access Point to limit public access, and restrict access by object prefixes, and object tags

Limit Access to VPC

You can create Access Points that limit all S3 storage access to a Virtual Private Cloud (VPC)

Easily scale your access

Access Points are easy to scale as you build more applications for your large shared data sets

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

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.

Use Amazon S3 VPC Endpoints and Service Control Policies

Amazon S3 VPC Endpoint Policies (Amazon S3 VPCEs)

For applications running in an Amazon Virtual Private Cloud (VPC), there are a few mechanisms to securely access Amazon S3. Amazon S3 VPCE policies allow you to access Amazon S3 over the AWS Private Network (instead of the internet) and attach a policy to limit which Amazon S3 resources can be accessed via your Amazon VPC.

They can also limit bucket access to specific VPCEs. On the receiving end, you can determine which external services can "talk to" the bucket in question. It is a best practice to include these parameters within your bucket policies.

AWS Organizations Service Control Policies

Using AWS Organizations, you can automate account creation, create groups of accounts to reflect your business needs, and apply policies for these groups for governance. AWS Organizations allows you to restrict what services and actions are allowed in your accounts.

With Service Control Policies (SCPs), you can apply permission guardrails on AWS IAM users and roles. For example, you can apply an SCP that restricts users in accounts in your organization from launching any resources in regions that you do not explicitly allow. Service Control Policies enable account-wide allow-list or deny-list of specific Amazon S3 API actions on buckets and objects within the account, regardless of the IAM identity of the requestor.

Ensure that your Amazon S3 buckets are not publicly accessible

Use Amazon S3 Block Public Access

Amazon S3 is the only object storage service that allows you to block public access to all of your objects at the bucket or the account level, now and in the future by using S3 Block Public Access.

To ensure that public access to all your new Amazon S3 buckets and objects is blocked, Amazon S3 Block Public Access is turned on by default for all new buckets. With a few clicks in the S3 management console, you can apply Amazon S3 Block Public Access to every bucket in your account – both existing and any new buckets created in the future.

AWS recommends that you turn on Amazon S3 Block Public Access for all 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 settings to suit your specific storage use cases.

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

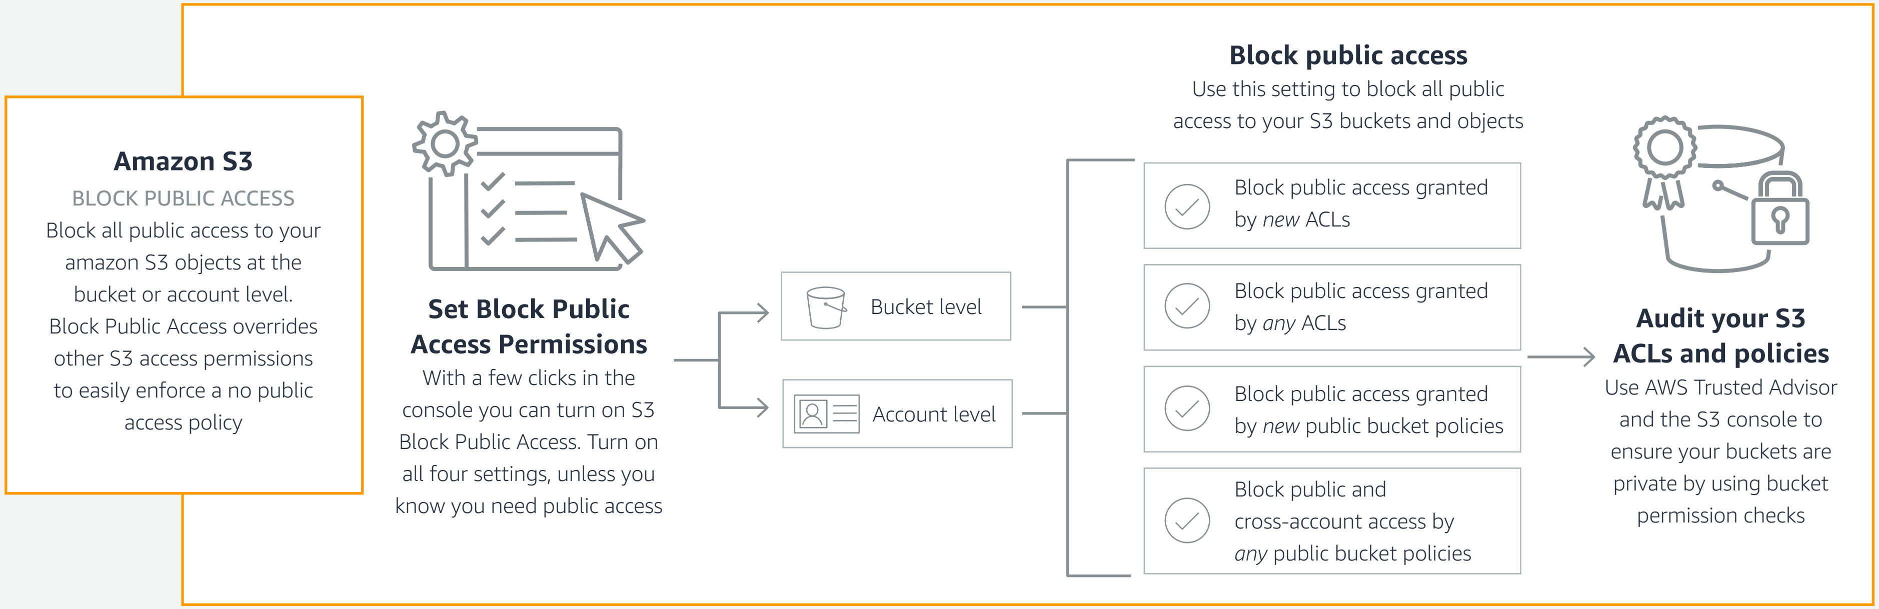

Amazon S3

BLOCK PUBLIC ACCESS

Block all public access to your amazon S3 objects at the bucket or account level. Block Public Access overrides other S3 access permissions to easily enforce a no public access policy

Set Block Public Access Permissions

With a few clicks in the console you can turn on S3 Block Public Access. Turn on all four settings, unless you know you need public access

Bucket level

Account level

Block public access

Use this setting to block all public access to your S3 buckets and objects

Block public access granted by *new* ACLs

Block public access granted by *any* ACLs

Block public access granted by *new* public bucket policies

Block public and cross-account access by *any* public bucket policies

Audit your S3 ACLs and policies

Use AWS Trusted Advisor and the S3 console to ensure your buckets are private by using bucket permission checks

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

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. 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, public downloads, 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

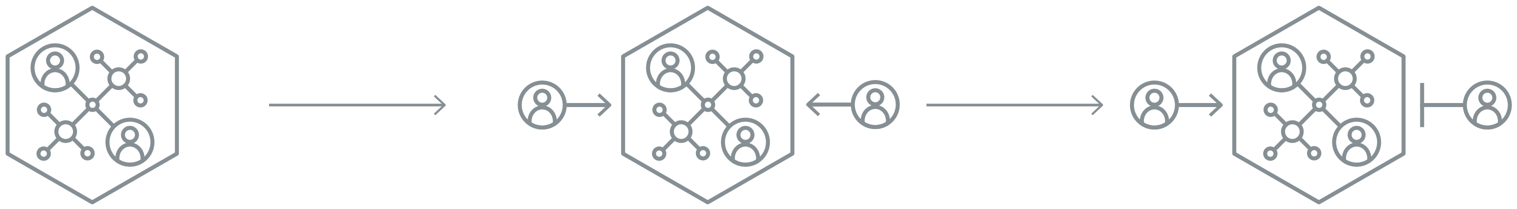

STEP 1 - Create an analyzer

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

Use AWS Trusted Advisor to check permissions and server logs

AWS Trusted Advisor checks Amazon S3 buckets that have open access permissions. Bucket permissions that grant List access to everyone can result in higher than expected charges if objects in the bucket are listed by unintended users at a high frequency. Bucket permissions that grant Upload/Delete access to everyone create potential security issues by allowing anyone to add, modify, or remove items in a bucket. This check examines explicit bucket permissions and associated bucket policies that might override the bucket permissions.

You can also use AWS Trusted Advisor to check the logging configuration of Amazon S3 buckets. Detailed access logs are delivered hourly to a bucket that you choose, and an access log record contains details about each request, such as the request type, the resources specified in the request, and the time and date the request was processed. 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.

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 intrinsic features and capabilities to provide the highest levels of data protection when it is used as the core storage for your applications.

Ensure that your Amazon S3 buckets are not publicly accessible

Data protection rests on the durability of the storage platform used. Durability is defined as the ability to protect data assets against corruption and loss. Amazon S3 is storage infrastructure for mission-critical and primary storage, and is designed to provide 99.999999999% data durability, which is 4 to 6 orders of magnitude greater than most on-premises, single-site storage platforms can provide.

Objects are redundantly stored on multiple devices across multiple facilities in an AWS Region. To help better ensure data durability, Amazon S3 operations synchronously store your data across multiple facilities, and is designed to sustain 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, that best practice includes secure access permissions, Cross-Region Replication, versioning, and a functioning, regularly tested backup.

Amazon S3 is designed to provide **99.999999999%** durability of objects over a given year. This durability level corresponds to an average annual expected loss of **0.000000001%** of objects. For example, if you store 10,000,000 objects with Amazon S3, you can on average expect to incur a loss of a single object once every 10,000 years.

Encrypt Everything

Turn on Amazon S3 Default Encryption

Amazon S3 default encryption provides a way to set the default encryption behavior for an S3 bucket. You can set default encryption on a bucket so that all new objects are encrypted when they are stored in the bucket. The objects are encrypted using server-side encryption with either Amazon S3-managed keys (SSE-S3) or customer master keys (CMKs) stored in AWS Key Management Service (AWS KMS).

It is important to note that when this is enabled, it does not encrypt existing objects within the bucket, and it does not force new or existing objects to be encrypted using the same encryption algorithms or keys.

Protecting 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.

You have three mutually exclusive options, depending on how you choose to manage the encryption keys.

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 master key that it regularly rotates. Amazon S3 server-side encryption uses one of the strongest block ciphers available, 256-bit Advanced Encryption Standard (AES-256), to encrypt your data. For more information, see Protecting Data Using Server-Side Encryption with SSE-S3.

2: Server-Side Encryption with Customer Master Keys (CMKs) Stored in AWS Key Management Service (SSE-KMS)

Server-Side Encryption with Customer Master Keys (CMKs) Stored in AWS Key Management Service (SSE-KMS) is similar to SSE-S3, but with some additional benefits and charges for using this service. There are separate permissions for the use of a CMK that provides added protection against unauthorized access of your objects in Amazon S3. SSE-KMS also provides you with an audit trail that shows when your CMK was used and by whom. For more information, see Protecting Data Using Server-Side Encryption with CMKs Stored in SSE-KMS.

3: 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, see Protecting data using server-side encryption with SSE-C.

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 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. You should allow only encrypted connections using HTTPS over TLS on Amazon S3 bucket policies.

Use versioning to preserve, retrieve, and restore your objects

Beyond core data protection, another key element 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.

Amazon S3 further protects your data using versioning. You can use versioning to preserve, retrieve, and restore every version of every object that is stored in your Amazon S3 bucket. With Amazon S3 Versioning, you can easily recover from both unintended user actions and application failures. Amazon S3 versioning can be used for data protection and retention scenarios such as recovering objects that have been accidentally deleted or overwritten and archiving previous versions of objects to Amazon S3 Glacier or Amazon S3 Glacier Deep Archive for long-term low-cost storage.

Trusted Advisor can also check for Amazon S3 buckets that do not have versioning enabled, or have versioning suspended. When versioning is enabled, you can easily recover from both unintended user actions and application failures.

Enable Multi-factor Authentication (MFA) Delete

You can use S3 Versioning Multi-Factor Authentication (MFA) Delete capability to provide an additional layer of security. By default, all requests to your Amazon S3 bucket require your AWS account credentials. If you enable Versioning with MFA Delete on your Amazon 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 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 Amazon 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 Versioning with MFA Delete, including how to purchase and activate an authentication device, please refer to the Amazon S3 Technical Documentation.

Amazon S3 Object Lock to enforce retention policies

Amazon S3 Object Lock blocks object version 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. You should use Amazon S3 Object Lock if you have regulatory requirements that specify that data must be WORM protected, or if you want to add an additional layer of protection to data in Amazon S3.

Used with S3 Versioning, which protects objects from being overwritten, you're able to ensure that objects remain immutable for as long as Amazon S3 Object Lock protection is applied.

You can configure S3 Object Lock in two Modes.

- Governance Mode: When deployed in Governance Mode, AWS accounts with specific IAM permissions are able to remove WORM protection from an object.
- Compliance Mode: If you require stronger immutability in order to comply with regulations, you can use Compliance Mode. In Compliance Mode, WORM protection cannot be removed by any user, including the root account.

Amazon S3 Replication to protect your data and meet compliance requirements

Amazon S3 is designed to provide 99.999999999% of 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 Amazon S3 buckets. With Amazon S3 Replication, you can configure Amazon S3 to automatically replicate Amazon S3 objects across different AWS Regions, or the region they currently reside in. All data assets are encrypted during transit with SSL to help achieve the highest levels of data security.

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 Amazon S3 Glacier, Amazon S3 Glacier Deep Archive 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

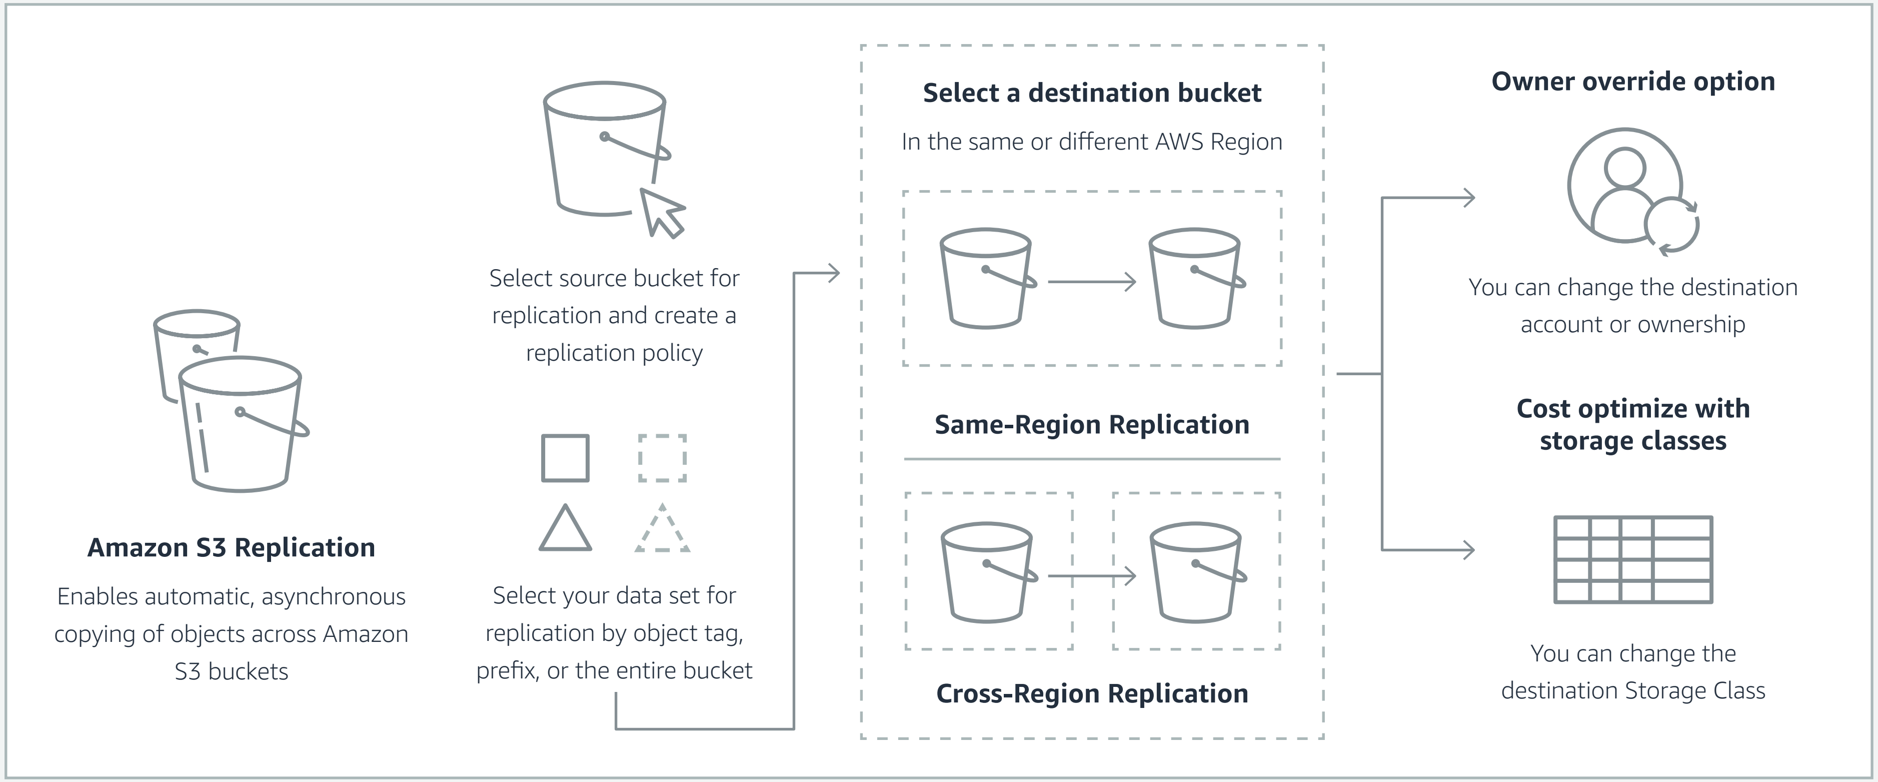

Amazon S3 Replication

Enables automatic, asynchronous copying of objects across Amazon S3 buckets

Select source bucket for replication and create a replication policy

Select your data set for replication by object tag, prefix, or the entire bucket

Select a destination bucket

In the same or different AWS Region

Same-Region Replication

Cross-Region Replication

Owner override option

You can change the destination account or ownership

Cost optimize with storage classes

You can change the destination Storage Class

Amazon S3 Replication Time Control

Amazon S3 replication Time Control helps you meet compliance “or business requirements” for data replication and provides visibility into Amazon S3 replication activity. Replication time control replicates most objects “that you upload” to Amazon S3 in seconds, and is backed by a Service Level Agreement (SLA) on the replication of 99.9% of objects within 15 minutes during any billing month

How Amazon S3 Replication Time Control works

Following these best practices will help 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.

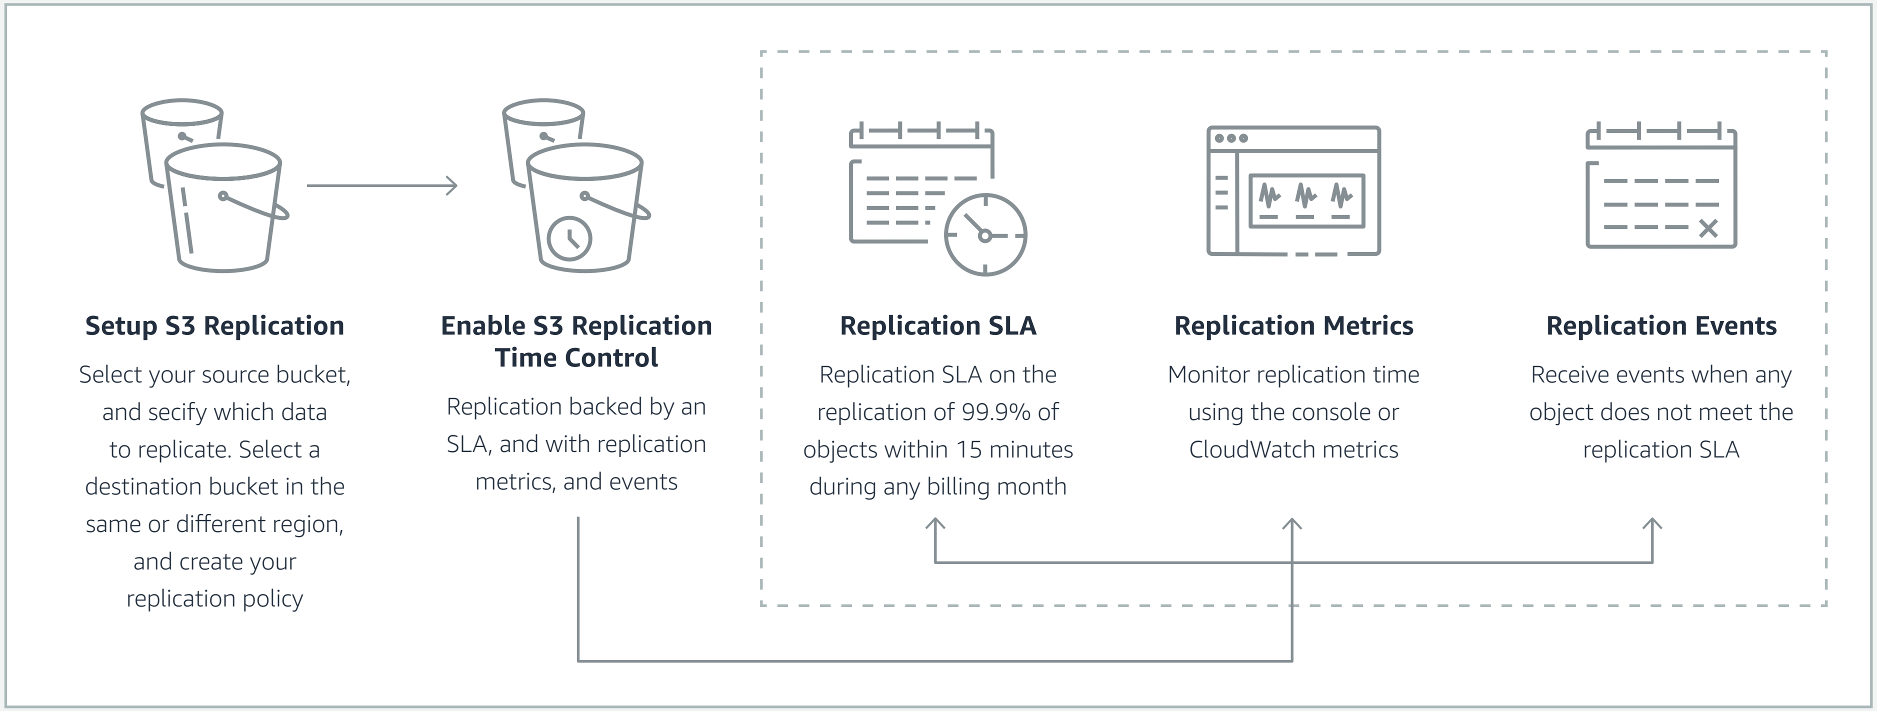

Setup S3 Replication

Select your source bucket, and seify which data to replicate. Select a destination bucket in the same or different region, and create your replication policy

Enable S3 Replication Time Control

Replication backed by an SLA, and with replication metrics, and events

Replication SLA

Replication SLA on the replication of 99.9% of objects within 15 minutes during any billing month

Replication Metrics

Monitor replication time using the console or CloudWatch metrics

Replication Events

Receive events when any object does not meet the replication SLA

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 profiles to monitor data access behavior, threat intelligence, 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.

How GuardDuty works

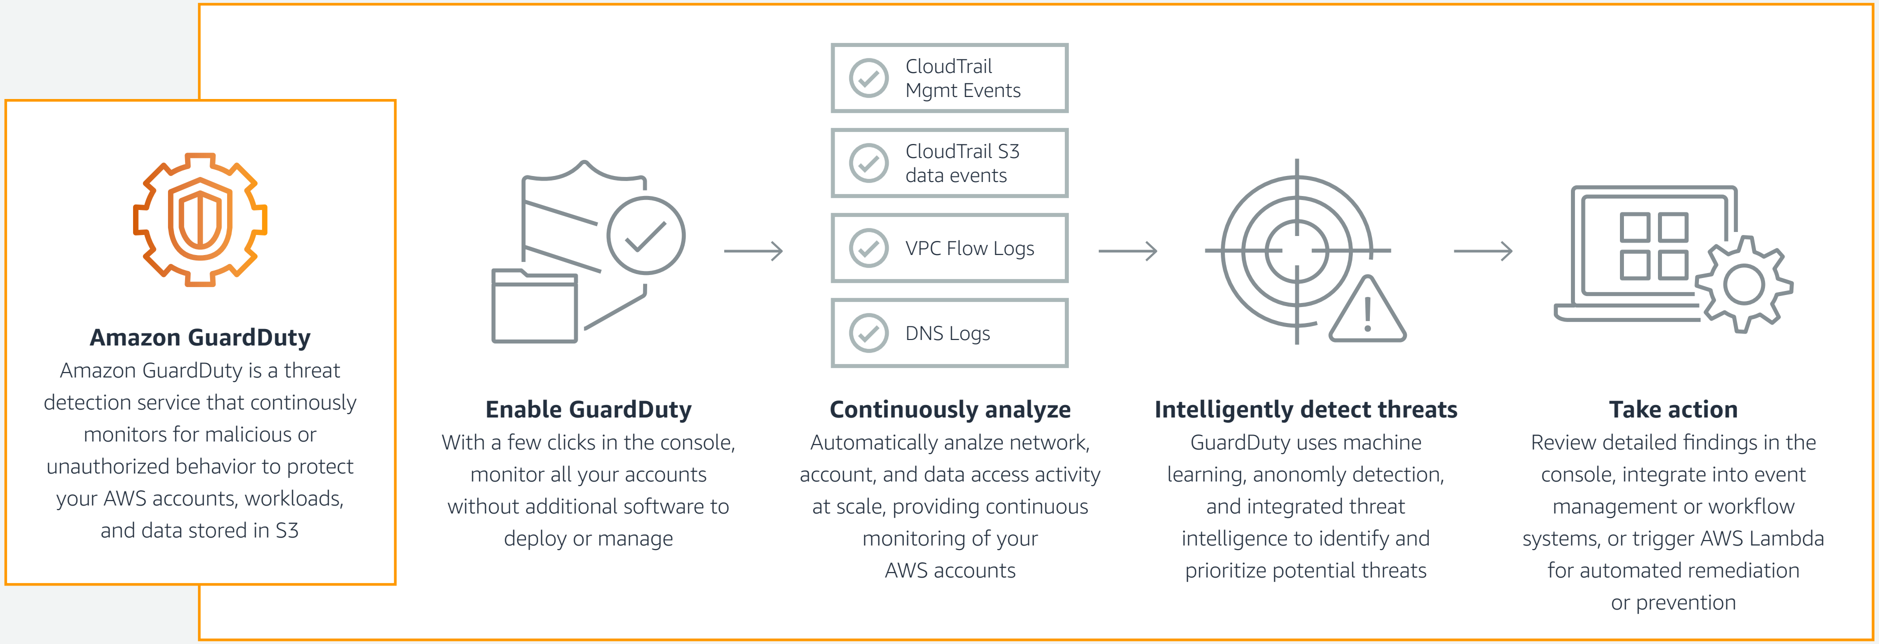

Amazon GuardDuty
Amazon GuardDuty is a threat detection service that continously monitors for malicious or unauthorized behavior to protect your AWS accounts, workloads, and data stored in S3

Enable GuardDuty
With a few clicks in the console, monitor all your accounts without additional software to deploy or manage

Continuously analyze
Automatically analze network, account, and data access activity at scale, providing continuous monitoring of your AWS accounts

Intelligently detect threats
GuardDuty uses machine learning, anonomly detection, and integrated threat intelligence to identify and prioritize potential threats

Take action
Review detailed findings in the console, integrate into event management or workflow systems, or trigger AWS Lambda for automated remediation or prevention

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.

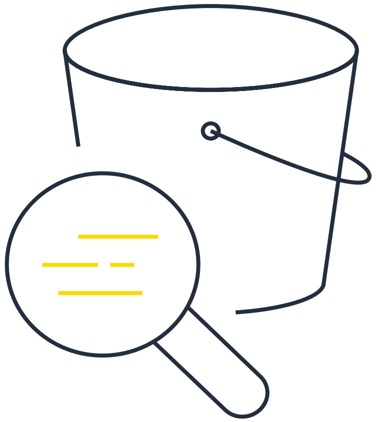

Use Tag Editor

to identify security-sensitive or audit-sensitive resources, then use those tags when you need to search for these resources.

Use Amazon S3 Inventory

to audit and report on the replication and encryption status of your objects for business, compliance, and regulatory needs.

Enable AWS Config

to assess, audit, and evaluate the configurations of your AWS and Amazon S3 resources. AWS Config monitors resource configurations, allowing you to evaluate the recorded configurations against the desired secure configurations. Using AWS Config, you can review changes in configurations and relationships between AWS resources, investigate detailed resource configuration histories, and determine your ability to meet your overalll compliance against the configurations specified in your internal guidelines. This can help you gain an understanding of your storage to give visbility to compliance auditing, security analysis, change management, and operational troubleshooting.

Inventory and audit using Amazon Macie,

as it automatically provides an inventory of Amazon S3 buckets including a list of unencrypted buckets, publicly accessible buckets, and buckets shared with AWS accounts outside those you have defined in AWS Organizations.

Implement monitoring of your S3 environment and bucket policies

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

Use Amazon Macie to gain visibility of your data security posture

Amazon Macie automates the discovery of sensitive data at scale and lowers the cost of protecting your data. Amazon Macie 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).

How Amazon Macie works

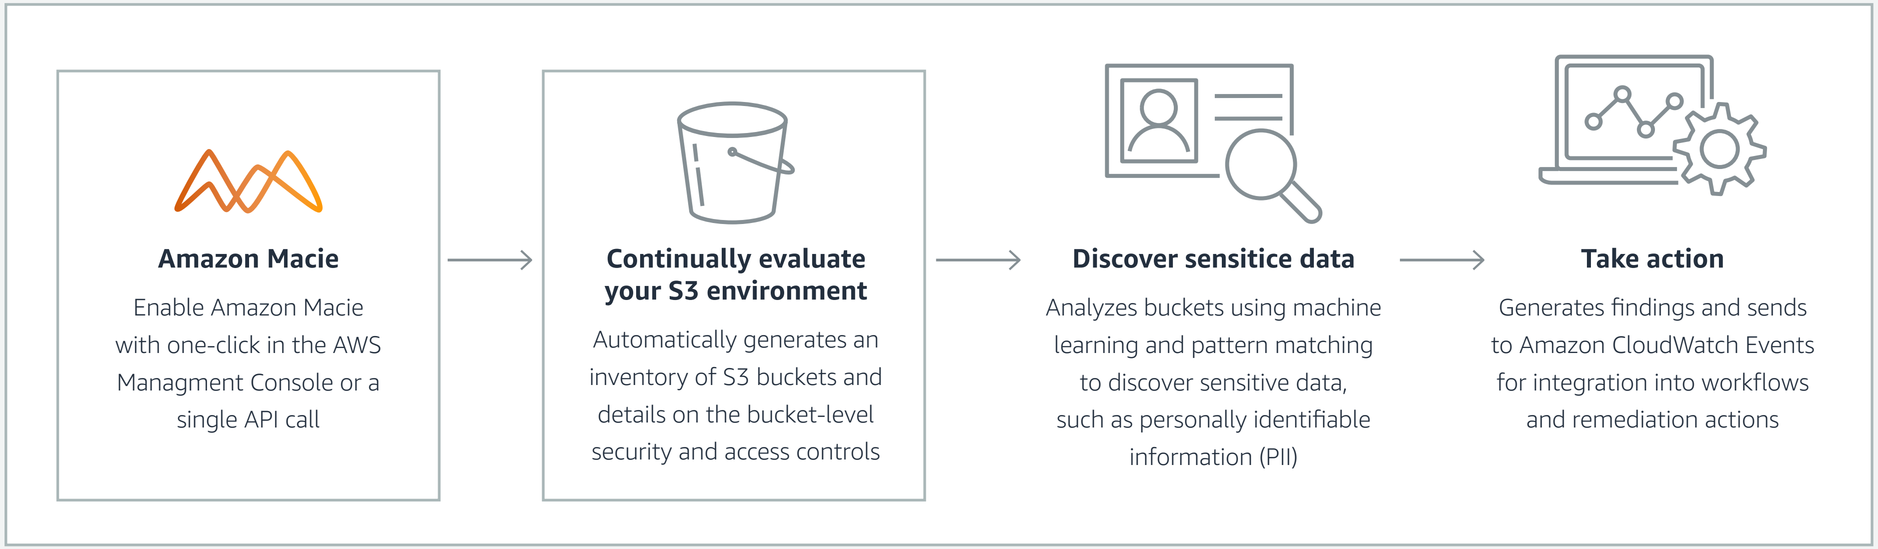

Amazon Macie

Enable Amazon Macie with one-click in the AWS Managment Console or a single API call

→

Continually evaluate your S3 environment

Automatically generates an inventory of S3 buckets and details on the bucket-level security and access controls

→

Discover sensitive data

Analyzes buckets using machine learning and pattern matching to discover sensitive data, such as personally identifiable information (PII)

→

Take action

Generates findings and sends to Amazon CloudWatch Events for integration into workflows and remediation actions

Use AWS CloudTrail to monitor and detect unusual activity

AWS CloudTrail provides a record of actions taken by a user, a role, or an AWS service in Amazon S3. You can use information collected by CloudTrail to determine the request that was made to Amazon S3, the IP address from which the request was made, who made the request, when it was made, and additional details.

When you set up your AWS account, CloudTrail is enabled by default. You can view recent events in the AWS CloudTrail console. To create an ongoing record of activity and events for your Amazon S3 buckets, you can create a trail in the CloudTrail console.

How AWS CloudTrail Works

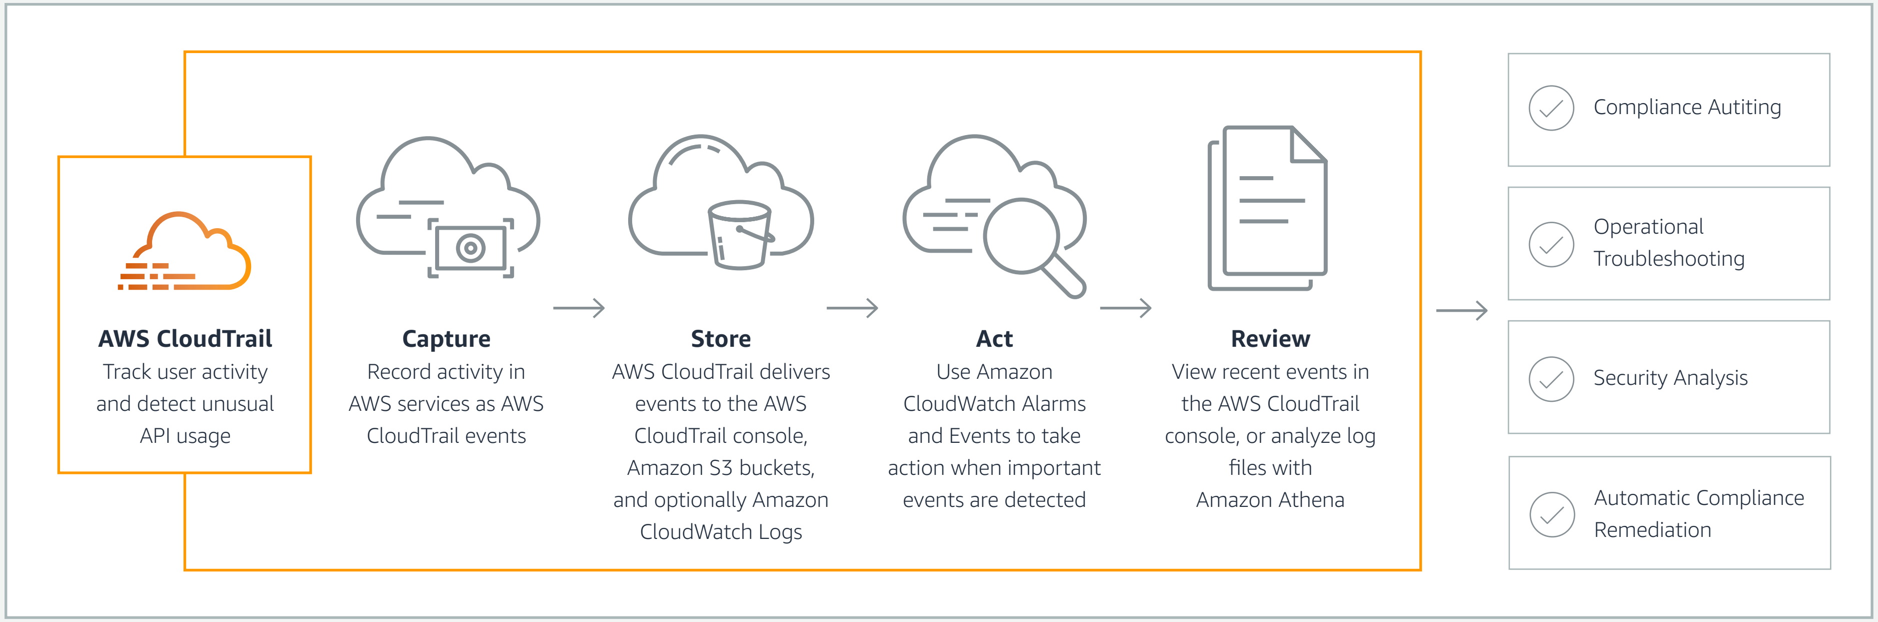

AWS CloudTrail
Track user activity and detect unusual API usage

Capture
Record activity in AWS services as AWS CloudTrail events

Store
AWS CloudTrail delivers events to the AWS CloudTrail console, Amazon S3 buckets, and optionally Amazon CloudWatch Logs

Act
Use Amazon CloudWatch Alarms and Events to take action when important events are detected

Review
View recent events in the AWS CloudTrail console, or analyze log files with Amazon Athena

- ✓ Compliance Autiting
- ✓ Operational Troubleshooting
- ✓ Security Analysis
- ✓ Automatic Compliance Remediation

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 - Monitor bucket storage using CloudWatch, which collects and processes storage data from Amazon S3 into readable, daily metrics. These storage metrics for Amazon S3 are reported once per day and are provided to all customers at no additional cost.
- Request metrics - Monitor Amazon S3 requests to quickly identify and act on operational issues.
- 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.

How Amazon CloudWatch works

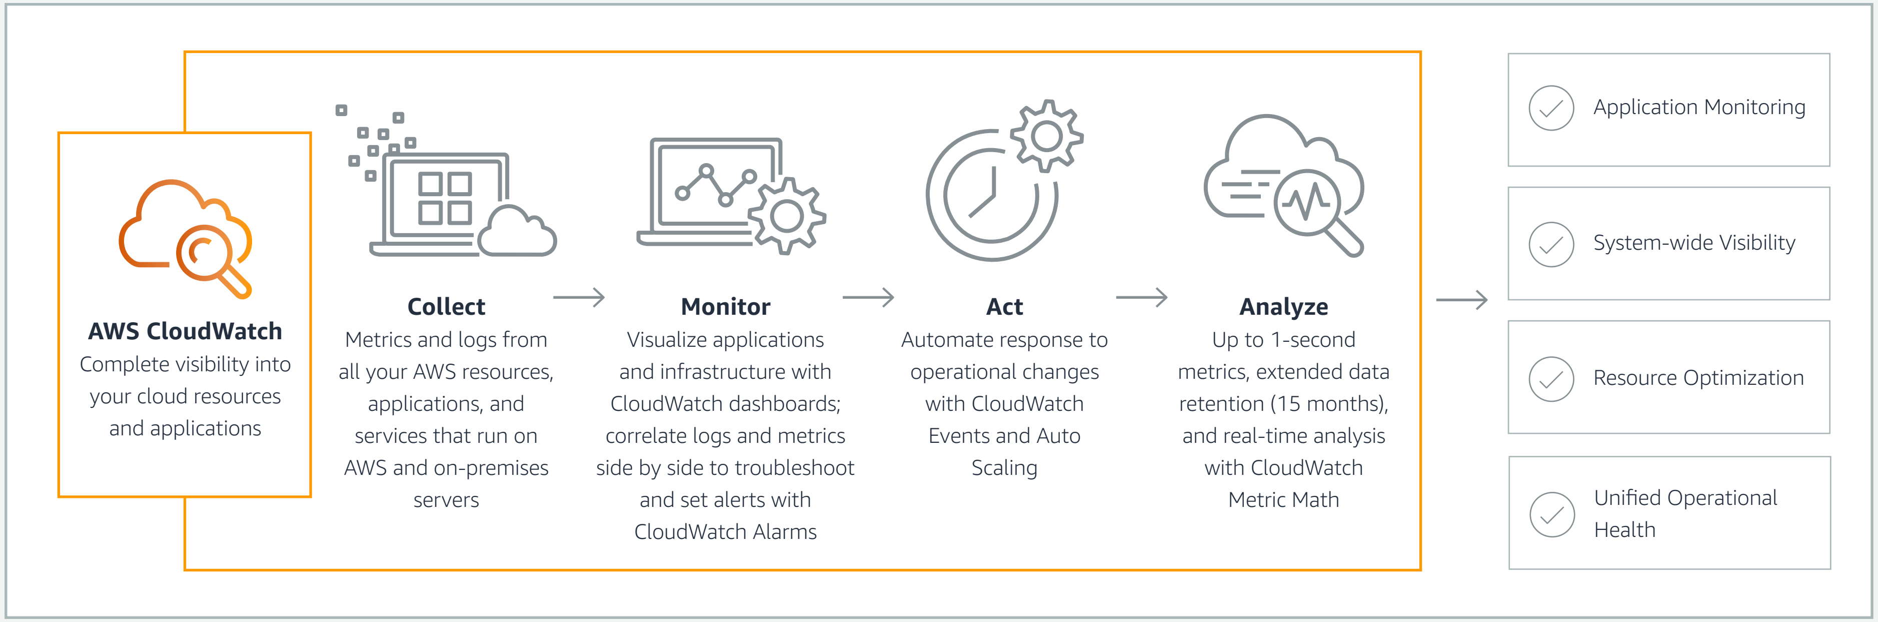

AWS CloudWatch
Complete visibility into your cloud resources and applications

Collect
Metrics and logs from all your AWS resources, applications, and services that run on AWS and on-premises servers

→

Monitor
Visualize applications and infrastructure with CloudWatch dashboards; correlate logs and metrics side by side to troubleshoot and set alerts with CloudWatch Alarms

→

Act
Automate response to operational changes with CloudWatch Events and Auto Scaling

→

Analyze
Up to 1-second metrics, extended data retention (15 months), and real-time analysis with CloudWatch Metric Math

→

✓ Application Monitoring

✓ System-wide Visibility

✓ Resource Optimization

✓ Unified Operational Health

Takeaways for Amazon S3 security best practices

Foundational tenets of securing your Amazon S3 buckets

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

Access Controls

- Incrementally grant permissions
- Leverage Amazon S3 bucket policies over bucket access control lists (ACLs)
- Enable account-level Block Public Access
- Leverage Access Points to scope application permissions
- Enable VPC endpoints with bucket policies limiting access
- Create and monitor IAM roles

Data Protection

- Encrypt everything
- Use bucket policy to enforce TLS
- Implement server and client-side encryption
- Enable Object Lock, versioning, MFA delete to help protect data

Monitoring and Management

- Continuously monitor which buckets do not meet your security settings
- 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

Closing page

Amazon S3 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/

Webinar: Best practices for Amazon S3 Security with S3 access management tools and S3 Block Public Access

https://pages.awscloud.com/Best-Practices-for-Amazon-S3-Security-with-S3-Access-Management-Tools-and-S3-Block-Public-Access_2019_0815-STG_OD.html

Documentation: Amazon S3 Security

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