Top 10 security best practices for securing backups in AWS

AWS Prescriptive Guidance

AWS Prescriptive Guidance: Top 10 security best practices for securing backups in AWS

Table of Contents

Top 10 security best practices for securing backups in AWS

Ibukun (IB) Oyewumi, Waleed Bedair, Raul Radu, and Rishi Singla, Amazon Web Services

Security is a shared responsibility between Amazon Web Services (AWS) and customers. Customers have asked for ways to secure their backups in AWS. Because security practices constantly evolve to mitigate new risks, it's important that you conduct regular risk assessments to determine the applicability of security controls, and implement multiple layers of controls to mitigate risks to your data.

This document will guide you through a curated list of the top 10 security best practices for securing your backup data and operations in AWS. While this guide focuses on backup data and operations using the AWS Backup service, these recommended security best practices can be used with other backup solutions, such as backup tools from the AWS Marketplace.

Best practices

When you secure backups in AWS, we recommend using the following approach:

Step 1. Implement a backup strategy

A comprehensive backup strategy is an essential part of an organization's data protection plan to withstand, recover from, and reduce any impact that might be sustained because of a security event. Create an extensive backup strategy that defines which data must be backed up, how often data must be backed up, and how backup and recovery tasks will be monitored.

When you develop a comprehensive strategy for backing up and restoring data, first identify interruptions that might occur, and their potential business impact.

Your objective is to build a recovery strategy that brings your workload back up or avoids downtime within the acceptable recovery objectives, Recovery Time Objective (RTO) and Recovery Point Objective (RPO). RTO is the acceptable delay between the interruption of service and restoration of service. RPO is the acceptable amount of time since the last data recovery point. Consider a granular backup strategy that includes all of the following:

- Continuous backup cadence
- Point-in-Time Recovery (PITR)
- File-level recovery

- Application data–level recovery
- Volume-level recovery
- Instance-level recovery

Ransomware

A well-designed backup strategy should include actions that can protect and recover your resources from ransomware, with detailed recovery requirements for your applications and their data dependencies. For example, while you establish preventive and detective controls to mitigate the risk of ransomware, make sure that you also design the appropriate level of granularity for cross-AWS Region or cross-account copy and restore patterns, to ensure that administrators do not restore corrupt backup data after the security event.

Retention requirements

In some industries, when developing a backup strategy, you must also consider the regulations for data retention requirements. Make sure that your backup strategy is designed with retention requirements that are sufficient to meet your regulatory needs for each data classification level and resource type.

Compliance

Consult your security compliance teams to validate whether your backup resources and operations should be included in or segmented from the scope of your compliance programs. Including backup and recovery as a critical part of your security program will help you understand where data is across your environment and appropriately define compliance scope.

For architectural best practices for designing and operating reliable, secure, efficient, and cost-effective workloads in the cloud, see Backup and recovery approaches using AWS and Reliability Pillar – AWS Well-Architected Framework.

Step 2. Incorporate backup in DR and the BCP

Disaster recovery (DR) is the process of preparing for, responding to, and recovering from a disaster. An important part of your resiliency strategy, the DR plan concerns how your workload responds when a disaster strikes. A disaster could be a technical failure, a human action, or a natural event.

The Business Continuity Plan (BCP) outlines how an organization intends to continue normal business operations during an unplanned disruption.

Your DR plan should be a subset of your organization's BCP. The BCP should also include AWS Backup procedures. For example, a security event that affects production data might require you to invoke a DR plan that uses AWS Backup to fail over to backup data from another AWS Region. Ensure that your employees are familiar with and have practiced using AWS Backup along with your organizational procedures, so that if disaster strikes, your organization can continue its normal operations with little or no service disruption. More information about using AWS Backup is provided in other sections in this guide.

Step 3. Automate backup operations

Your organization's backup plans and resource assignments should be configured to reflect your enterprise data protection policies. By automating and deploying backup policies or organization-wide backup plans, you can standardize and scale your backup strategy. You can use AWS Organizations to centrally automate backup policies to implement, configure, manage, and govern backup activity across supported AWS Backup resources by scheduling backup operations.

To improve productivity and govern infrastructure operations across multiple-account environments, consider implementing infrastructure as code (IaC) and event-driven architecture as essential parts of your digital transformation and backup strategy. Automating backups provides the following benefits:

- Reduces manual overhead from time-consuming configuration of your backups
- Minimizes the risk for errors
- Provides visibility on drift detection
- Enhances backup policy compliance across multiple AWS workloads or accounts

Implementing backup policies as code can help you meet data protection regulations by doing the following:

- Configuring different requirements for your resource types
- Scaling your enterprise data protection strategy
- Implementing lifecycle rules to specify how long before a recovery point either transitions to cold storage or is deleted, which can help optimize your costs

When automating your backup operations, you can scale resource assignment options by using AWS tags and resource IDs to automatically identify the AWS resources that store data for your business-critical applications and protect your data using immutable backups. This can help you prioritize security controls, such as access permissions and backup plans or policies.

Step 4. Implement access control mechanisms

When thinking about security in the cloud, your foundational strategy should begin with a strong identity foundation to ensure a user has the right permissions to access data. Appropriate authentication and authorization can mitigate the risk of security events. The shared responsibility model requires AWS customers to implement access control policies. To create and manage access policies at scale, you can use AWS Identity and Access Management (IAM).

When configuring access rights and permissions, implement the principle of least privilege by ensuring each user or system accessing your backup data or vault is given only the permissions necessary to fulfill their job duties. Use AWS Backup to set access policies on backup vaults to protect your cloud workloads.

For example, by implementing access control policies, you can grant users access to create backup plans and on-demand backups while limiting their ability to delete recovery points. Using vault access policies, you can share a destination backup vault with a source AWS account, user, or IAM role, as required by your business needs. You can also use access policies to share a backup vault with one or multiple accounts, or with your entire organization in AWS Organizations. For more information, see the AWS Backup documentation.

As you scale your workloads or migrate into AWS, you might need to centrally manage permissions to your backup vaults and operations. Use service control policies (SCPs) to implement centralized control over the maximum available permissions for all accounts in your organization. SCPs offer defense in depth, and they ensure that your users stay within the defined access control guidelines. For more information, see Managing access to backups using service control policies with AWS Backup.

To mitigate security risks such as unintended access to your backup resources and data, use IAM Access Analyzer to identify any AWS Backup IAM role shared with the following:

- An external entity such as an AWS account
- A root user
- An IAM user or role

- A federated user
- An AWS service
- An anonymous user
- Any other entity that could be used to create a filter

Step 5. Encrypt backup data and vault

Organizations increasingly need to improve their data security strategy, and they might be required to meet data protection regulations as they scale in the cloud. The correct implementation of encryption methods can provide an additional layer of protection above foundational access control mechanisms. This added layer provides a mitigation if your primary access control policies fail.

For example, if you configure overly permissive access control policies on your AWS Backup data, your key management system or process can mitigate the maximum impact of a security event. This is because there are separate authorization mechanisms to access your data and encryption key, which means that the backup data is viewable only as cipher text.

To get the most from AWS Cloud encryption, encrypt data both in transit and at rest. To protect data in transit, AWS uses published API calls to access AWS Backup through the network using the TLS protocol to provide encryption between you, your application, and the AWS Backup service. To protect data at rest, you can use AWS cloud-native AWS Key Management Service (KMS) or AWS CloudHSM. A cloud-based hardware security model (HSM), AWS CloudHSM uses Advanced Encryption Standard (AES) with 256-bit keys (AES-256), a strong industry-adopted algorithm for encrypting data. Evaluate your data governance and regulatory requirements, and select the appropriate encryption service to encrypt your cloud data and backup vaults.

Encryption configuration differs depending on the resource type and backup operations across accounts or Regions. Certain resource types support the ability to encrypt your backups using a separate encryption key from the key used to encrypt the source resource. Because you are responsible for managing access controls to determine who can access your AWS Backup data or vault encryption keys and under which conditions, use the policy language offered by AWS KMS to define access controls on keys. You can also use AWS Backup Audit Manager to confirm that your backup is properly encrypted. For more information, see Encryption for backups in AWS Backup.

You can use AWS KMS multi-Region keys to replicate keys from one Region into another. Multi-Region keys are designed to simplify encryption management when your encrypted data has to be

copied into other Regions for disaster recovery. Evaluate the need to implement multi-Region AWS KMS keys as part of your overall backup strategy.

Step 6. Safeguard backups using immutable storage

Organizations can use immutable storage to write data in a Write Once Read Many (WORM) state. While in a WORM state, data can be written one time, and read and used as often as needed after it has been committed or written to the storage medium. Immutable storage helps to ensure that data integrity is maintained. It also provides protection against the following:

- Deletes
- Overwrites
- Inadvertent and unauthorized access
- Ransomware compromise

Immutable storage offers an efficient mechanism to address potential security events that might have real impacts on your business operations.

You can use immutable storage for better governance when paired with strong SCP restrictions. You can also use immutable storage in a compliance WORM mode when the letter of the law (such as a legal hold) requires access to immutable data.

To maintain data availability and integrity, you can use AWS Backup Vault Lock to protect your backups. When you use AWS Backup Vault Lock, unauthorized entities can't erase, alter, or corrupt your customer or business data during the required retention period. AWS Backup Vault Lock helps you meet your organization's data protection policies by preventing the following:

- Deletions by privileged users (including the AWS account root user)
- Changes to your backup lifecycle settings
- Updates that alter your defined retention period

AWS Backup Vault Lock ensures immutability and adds an additional layer of defense that protects backups (recovery points) in your backup vaults. This is especially useful in highly regulated industries that have stringent integrity needs for backups and archives. AWS Backup Vault Lock makes sure your data is preserved along with a backup to recover from in case of unintended or malicious actions.

Important

- AWS Backup Vault Lock has not yet been assessed for compliance with the Securities and Exchange Commission (SEC) rule 17a-4(f) and the Commodity Futures Trading Commission (CFTC) in regulation 17 C.F.R. 1.31(b)-(c).
- AWS Backup Vault Lock takes effect immediately. It gives you a *minimum three-day (72-hour) cooling-off period* to delete or update its configuration before it permanently locks your vault. You can optionally extend the duration of this cooling-off period. Use this cooling-off period to test AWS Backup Vault Lock against your workloads and use cases. After your cooling-off period expires, neither you nor AWS Support can delete or otherwise alter AWS Backup Vault Lock or the contents of your locked vault. For more information, see the documentation on AWS Backup Vault Lock.

Step 7. Implement backup monitoring and alerting

Backup jobs can fail. A failed job, such as backup, restore, or copy task, might impact on subsequent steps in a process. When the initial backup job fails, there's a high probability that other succeeding tasks will also fail. In such a scenario, you can best understand the course of events through monitoring and notification. Monitoring and alerting can provide organizational awareness for your backup jobs, which helps you respond to backup failures.

Activating and configuring notifications to monitor AWS Backup jobs provides the following benefits:

- Gives you awareness of your backup activities
- Helps ensure that you meet critical service-level agreements (SLAs)
- Enhances your business-as-usual monitoring
- Helps you meet compliance obligations

You can implement backup monitoring for your workloads by integrating AWS Backup with other AWS services and ticketing systems to perform automated investigation and escalation flows. For example, you can do the following:

- Use Amazon CloudWatch to track metrics, create alarms, and view dashboards.
- Use Amazon EventBridge to monitor AWS Backup processes and events.
- Use AWS CloudTrail to monitor AWS Backup API calls with detailed information on the time, source IP, users, and accounts making those calls.
- Use Amazon Simple Notification Service (Amazon SNS) to subscribe to AWS Backup related topics such as backup, restore, and copy events.

You can use AWS Backup Audit Manager to automatically generate evidence of your daily backup audit reports for each account and Region. You can also scale your backup monitoring across multiple accounts by using a set of automation templates and dashboards (known as the backup observer solution) to obtain aggregated daily cross-account multi-Region AWS Backup reporting.

Step 8. Audit backup configuration

To ensure that the backup program is performing as it should and to identify and correct any anomalies from backup processes, audit the compliance of AWS Backup policies against defined controls such as defined backup frequency. To find and investigate backup operations or resources that are not compliant with your business requirements, continuously and automatically track your backup activity and generate automatic reports.

AWS Backup Audit Manager provides built-in, customizable compliance controls that align with your business compliance and regulatory requirements. You can use prebuilt and customizable controls as audit frameworks to evaluate your AWS Backup practices. The controls include:

- Backup resources protected by backup plans
- Backup plan minimum frequency and minimum retention
- Backup recovery point encrypted
- Backup recovery point manual deletion
- Backup recovery point minimum retention
- Cross-Region copy
- Cross-account copy
- Backup Vault Lock

For infrastructure as-code (IaC) automation, you can use AWS Backup Audit Manager with AWS CloudFormation.

AWS Security Hub provides you with a comprehensive view of your security state in AWS. It also helps you check your environment against security best practices and industry standards such as AWS Foundational Security Best Practices controls. If you use Security Hub within your cloud environment, we recommend you enable the AWS Foundational Security Best Practices standard, because it includes detective controls that can help with securing backups in AWS. Most of the detective controls in AWS Backup Audit Manager and Security Hub are also available as AWS Config managed rules.

Step 9. Test data recovery capabilities

Your backup strategy must include testing your backups. A backup strategy is not effective if backed up data cannot be restored. Regularly test your ability to find certain recovery points and restore them.

While AWS Backup automatically copies tags from the resources it protects to the recovery points, tags are not by default copied from recovery points to the corresponding restored resources. To scale your inventory management and locate recovery points, consider using AWS Backup events to initiate a tag replication process on resources created by AWS Backup restore jobs.

You can start your data recovery workflow by establishing data recovery patterns and then regularly testing them. To increase confidence in your ability to recover backup data, create a basic, repeatable process for performing continuous data recovery testing. For example, you can create a pattern to test a cross-account, cross-Region restore operation from a central DR backup vault encrypted with a customer-managed AWS KMS key to a source account backup vault encrypted with a different customer-managed AWS KMS key.

If you don't frequently test such restore operations, you might find that your assumptions on AWS KMS encryption for cross-account, cross-Region operations are incorrect. Often, the only backup recovery pattern that actually works is the path you test frequently. Through routine testing of supported backup resource types, you can spot early warnings that could potentially cause future disturbances and loss of critical data. If possible, maintain a limited but feasible number of recovery paths and patterns to prevent wasted storage space, optimize costs, and save time. Fixing the problem when a recovery test fails is easier than losing valuable or critical data.

Step 10. Incorporate backup in your incident response plan

Security Incident Response Simulations (SIRS) are internal events that provide a structured opportunity to practice your incident response plan and procedures during a realistic scenario.

It's valuable to test your backup data and operations in creative SIRS activities to test yourself against the unexpected. This helps you validate your organizational readiness and develop comfort with the rare and unexpected. Your simulations must be realistic, and they should involve cross-functional organizational teams that are required to respond to events.

Start with basic simulation exercises, and work toward a full, complex event. For example, you can build a realistic model that consists of a virtual private cloud (VPC) and associated resources that simulate inadvertent overexposure of information or a potential data breach caused by changes to policies and access control lists. To evaluate how well your incident response plan worked, document lessons learned, and identify improvements that need to be made to future response procedures.

You can use AWS Backup to set up automated instance-level backups as Amazon Machine Images (AMIs) and volume-level backups as snapshots across multiple AWS accounts. This can help your incident response team enhance their forensic processes, such as automated forensic disk collection, by providing a restore point that could reduce the scope and impact of potential security events such as ransomware.

Next steps

This guide covered the top 10 security best practices and controls to protect your backup data in AWS. We recommend using these best practices to design and implement a backup and recovery strategy and architecture, with multiple layers of controls, that scale and achieve your business needs. For more information about AWS Backup, see the AWS Backup documentation.

If you have questions about this guide, start a new thread on the AWS Backup forum or contact contact AWS Support.

Resources

AWS Prescriptive Guidance

- Backup and recovery approaches on AWS (guide)

AWS documentation

- AWS Backup
- AWS CloudFormation
- AWS CloudHSM
- AWS CloudTrail
- Amazon CloudWatch
- AWS Config
- Amazon EventBridge
- IAM
- AWS KMS
- AWS Organizations
- AWS Security Hub
- Amazon SNS

Blog posts

- Automate centralized backup at scale across AWS services using AWS Backup
- Disaster Recovery (DR) Architecture on AWS, Part I: Strategies for Recovery in the Cloud
- The importance of encryption and how AWS can help
- Enhance the security posture of your backups with AWS Backup Vault Lock
- Monitor, Evaluate, and Demonstrate Backup Compliance with AWS Backup Audit Manager
- Create and share encrypted backups across accounts and Regions using AWS Backup
- Simplify auditing your data protection policies with AWS Backup Audit Manager
- Managing access to backups using service control policies with AWS Backup

- Obtain aggregated daily cross-account multi-Region AWS Backup reporting

GitHub repositories

The following code repositories provide you with a management and deployment solution for implementing backup and recovery with AWS Backup across your AWS Organizations organization spanning multiple accounts and AWS Regions.

- Implement AWS Backup using AWS
- Implement AWS Backup using CI/CD Pipelines

Document history

The following table describes significant changes to this guide.

Change	Description	Date
Initial publication	—	May 13, 2022