



10 Best Practices for AWS

A security, backup and integration checklist
for success with AWS



Table of contents

Introduction	3
Evaluation	3
Outline Your Data Protection Strategy	4
Deployment and Configuration	5
Security	5
Data and Application Classification	6
Compliance and Archive	7
Operate, Restore and Optimize	8
Consider a Policy-Based Approach for Total Coverage	8
Restore Methodology	9
Metering, Cost, and Fees	10
Monitoring, Reporting, and Testing	11
Evolve Beyond Backup	12
Day 2 Operations	12
Plan for the Future	13
Summary	14

Introduction

Whether you are a seasoned pro on Amazon Web Services (AWS), or just getting started, you likely know your data on AWS is just that — YOUR data. And that data isn't going to protect itself... Every organization using AWS should have a robust and secure backup and recovery solution that ensures data resilience and fits their needs.

Irrespective of what drives your needs — security, compliance, speed of recovery, and more — these are 10 best practices everyone should follow to help you optimize your costs, secure your data, and free up your time.



Evaluation

Outline Your Data Protection Strategy

To design a backup and recovery strategy that fits your organization's needs, you need to answer these questions to help you define various aspects of your data protection and which tools can meet your requirements.

At a high level, it is best to align on the following:

- **How much data loss can your organization accept?**
Referred to as Recovery Point Objective or RPO, answering this question will help you understand how often an application or data needs to be protected, e.g., continuously, hourly, daily, etc.
- **How quickly after an outage does an application need to be available?**
Referred to as Recovery Time Objective or RTO, this measurement not only defines the methods of recovery you need to meet requirements (e.g., full or file-level recovery), but also how you protect and store the data, (e.g., snapshots, backups and replicas).
- **How long does backup data need to be retained?**
Retention is often dictated by corporate mandates and/or regulatory compliance regarding when, where and how long data can be stored for.

Note that these variables are rarely the same across all workloads in your environment. Rather, they vary depending on the criticality of the application and data, as well as the governance and retention of data.

To help meet these objectives, you typically have a few options:



Snapshot

Capture data at a point-in-time which is immediately available for applications that require read-only data.



Backup

Create copies of the data set and independently store and/or archive it.



Replication

Replicate data across infrastructure, often used for mission-critical applications that have little tolerance for downtime.



Architecture

Design and implement a high availability architecture, including Multi-Availability Zones (Multi-AZ), clustering, auto scaling and load balancing to complement the above.



Veeam provides versatile data protection solutions that seamlessly integrate with any design of the AWS environment, enabling organizations to efficiently manage and protect their data regardless of architecture. Whether organizations adopt a single-cloud or hybrid cloud approach, Veeam offers comprehensive backup, recovery, and replication capabilities that are compatible with various AWS services and deployment models. This flexibility allows organizations to design their AWS environment based on unique requirements while ensuring critical data is always protected, accessible, and recoverable.

Deployment and Configuration

Security

With cyberthreats being at an all-time high, having a backup and recovery plan is essential to any effective cybersecurity strategy. To ensure the security and dependability of backups, there are several best practices you must implement during deployment related to security.

For instance, relying on a single account for production and protection could result in everything becoming compromised in the event of a successful attack, leaving you without recovery options. In AWS, the best practice is to isolate your backup data from the production environment. This logical separation can be accomplished by conducting backups across separate accounts.

Along with logically separating your data, you should also utilize the security-focused AWS capabilities to further protect your data against cyberattacks and human error. Some examples include:



Access Management

Granularly control who has access to specific data, applications, services and functions. This can be achieved with fine-grain management, rotation and deletion of AWS Identity and Access Management (IAM) roles, as well as third-party options like multi-factor authentication and role-based access control (RBAC).



Immutability

Putting backups in a write once read many (WORM) state ensures your data can't be changed, encrypted or deleted, even if it gets hit by a ransomware attack. Amazon S3 Object Lock, an immutable data security option, adds an additional layer of protection to your backup data.



Encryption

Encrypting data in-flight and at rest can prevent unauthorized access and exfiltration of backups. Encryption options include AWS Key Management Service (AWS KMS) — particularly Customer Managed Keys (CMKs) for greater control — as well as encryption offerings from that backup vendor.



NOTE:

Align immutability policies with retention requirements, otherwise you run the risk of data not being appropriately deleted and subsequently increasing cost and compromising compliance with unnecessary lengthy retention.

Data and Application Classification

Before creating policies or jobs to secure your applications and data hosted on AWS, it's important to classify them. This will ensure that the appropriate level of protection and retention is applied. The classification will probably be based on the concepts of RPOs, RTOs, and retention discussed previously in this guide.

One of the most effective strategies for classification (and further down the line, automated protection) is to define and apply tags to your AWS resources through the AWS Management Console. These tags should be aligned with specific policy settings, ensuring all current and future workloads are appropriately protected and compliant without overspending. An example of tagging could use nomenclature like gold, silver and bronze, whereby:

- **Gold**
Reserved for the most critical applications and data that demand the lowest RTOs and RPOs and requires long-term retention across multiple storage tiers.
- **Silver**
Utilized for critical applications and data that require low RTOs and RPOs, but have shorter retention periods.
- **Bronze**
For applications and data sets that can tolerate greater data loss and longer periods of downtime, with short retention periods.



AWS allows for up to 50 tags per resource that you can define and apply, allowing for greater flexibility in terms of protection and retention standards while still benefiting from automation to free up time and resources.

Compliance and Archive

Every organization is required to adhere to regulatory and corporate standards for their data, whether highly sensitive or not. These requirements can include the RPOs, RTOs and retention, but can also extend to when it must be deleted, the type of media it must be stored on, as well as its location (e.g., state, country, region, etc.). Therefore, make sure you understand and stay up to date with governing data protection policies and regulatory requirements. Some of the necessary questions you should ask yourself include:

- **Are there laws and regulations that govern how my company protects data? For example:**
 - General Data Protection Regulation (GDPR)
 - Sarbanes-Oxley Act (SOX)
 - Payment Card Industry (PCI)
 - Health Insurance Portability and Accountability Act (HIPAA)
- **Are there individuals in my organization that can help me identify regulations and achieve compliance?**
- **Does my platform and/or infrastructure meet these requirements (e.g., media type, physical location, accessibility, etc.)?**
- **When is it safe and appropriate to delete data?**

After comprehending the compliance requirements, you can begin to appropriately design and automate policies to correctly manage the lifecycle of your data. Demonstrating compliance through reporting is also critical so that you operate in a state that is always audit-ready.




Veeam's data protection solutions help enforce security measures such as encryption and access controls, ensuring the confidentiality and integrity of data within AWS environments. Veeam's data protection solutions coupled with AWS' security features can further enhance the overall security posture of your organization's environment.


In terms of classification, compliance, and archive, Veeam enables organizations to classify and categorize data based on its importance, allowing for more granular control over backup, retention, and archival processes. This classification capability aligns with compliance regulations and helps organizations meet data residency, privacy, and retention requirements. Combined with the secure and durable storage options provided by AWS, organizations have the flexibility to securely store and retrieve archived data while adhering to governance, legal, regulatory, and compliance obligations.


Operate, Restore and Optimize


Consider a Policy-Based Approach for Total Coverage

With a solid foundation in place, it's time to implement your strategy. A policy-based approach should always be top of your list due to the many benefits that automation brings, including:

 **Optimizing cost**
Manual practices often lead to snapshot sprawl and excessive billing for data protection. By automating the lifecycle of backup data across storage classes and tiers, data can be retained appropriately without the subsequent overspend.


 **Eliminating unprotected data**
The challenges of shadow IT, or data housed on non-centralized systems, are exacerbated by the ease of creating new workloads in the cloud. Automating backup helps to ensure any new workload, application or data set deployed in the cloud will be automatically protected, especially when utilizing a tag, account, or region-based approach.

 **Achieving compliance**
Meeting compliance requirements for data protection and security is one thing; reporting on compliance it is another. Automation helps ensure that all data is appropriately accounted for and provides the benefits of monitoring, alerting, and reporting.

 **No more babysitting backup**
Automating previously scripted operations ensures data protection and security are implemented while freeing up personnel time and resources for other strategic priorities.



Policies should be based on the recommendations made earlier in this guide, dictated by the varying service level agreements (SLAs) and compliance requirements. Once defined, backup should ideally be a set-it-and-forget it operation.



Restore methodology

You are protecting your data for a reason — to recover it when lost or compromised. Your need to be able to restore your data as quickly and efficiently as possible — whether by instances, database or file share.


The power of choice is your friend here, so be sure to select a solution that provides you with all the necessary tools to tackle any scenario. What about the case when a single file is in need of recovery vs. an entire instance?

What if the source account has been compromised by ransomware? The ability to restore what, when, how, and where you need it, puts you in the seat of control and better assists you to meet RTOs.

Here are some of the available options for you:

- **Full restore** of an entire instance, database or file share
- **File-level restore**, including easy eDiscovery
- **Failover/failback** of replicated apps and data
- **In-place** of the existing file, folder or instance
- **As new**, including to a new account, region or platform



It is also critical to be mindful of just what is required from you when performing a restore.  Manual recoveries require you to create, connect, boot and mount instances and volumes which is time-consuming and error prone. Solutions that automate the entire recovery process will make you more effective at the time you need it most.

Metering, Cost, and Fees

Storing and recovering data in the cloud can get expensive, especially when considering the pay-for-what-you-use model for compute, storage, networking, API calls and more. Help prevent bill shock by understanding the total cost of the data protection and security operations of your AWS environment.

There are many resources, tools and best practices you can follow that help stop sprawling cloud costs without sacrificing the degree of protection.

- Storage classes and associated costs:** Each storage class varies in price depending on factors such as how much data you store, where you store it, and how frequently you access it. Current pricing can be confirmed with AWS, however:

Storage Type	Price
EBS Snapshots — Month	\$\$\$\$
S3 Standard — First 50/Month	\$\$\$
S3 Glacier Instant Retrieval — All Storage/Month	\$\$
S3 Glacier Deep Archive — All Storage/Month	\$

- Proactive calculators:** Allow you to get a close estimate for AWS monthly cost based on the list prices and policy parameters input by the user.
- Dynamic compute:** Keeps cost low by automatically instantiating and terminating data movers. Backup operations require compute, but backup operations are seldom 24/7/365.
- Cross-regional traffic:** Cross-region traffic incurs additional expense. By defining which services require replication across regions, typically by their mission criticality, helps mitigate overspend.



You can reduce the risk of unexpected bills by taking a proactive approach and designing an infrastructure that incorporates cost considerations.



Monitoring, Reporting, and Testing

In this perpetually evolving world, threats and risks can take new forms so, it is imperative for your AWS environment to be current. To have this level of assurance, you need to receive regular updates and insights without having to touch the tool. This may include:

- **Monitoring**
With cloud monitoring software you will be able to watch for irregularities in your instances, databases, storage, and data movement.
- **Alerting**
By sending routine or ad-hoc alerts to your email you can be notified when an incident arises allowing you to stay on top of activity.
- **Reporting**
When the above is set up properly you will be compliant internally and externally and be able to run reports automatically.

Even better, some tooling delivers API capabilities so that you can integrate the status of backup and recovery into other third-party applications that your organization utilizes.

To complete a thorough health check, continuous testing is essential. By being proactive, you may find a weak point in your security or catch a potential issue before it becomes destructive, like an unrecoverable backup. Do not be conservative with your testing. Cover a wide range of scenarios, from power outages caused by natural disasters to accidental file deletion to ransomware attacks. This is critical not only to validate the technical success, but also to practice and build muscle memory from a human perspective so that the first time a real-world recovery scenario arises, people are prepared and ready.



Veeam facilitates a policy-based approach by enabling organizations to define and enforce backup and recovery policies, ensuring consistency and reducing errors. With granular restore capabilities, users can quickly recover specific files or applications, minimizing downtime and increasing productivity. Veeam helps manage costs by optimizing storage utilization and offers monitoring, reporting, and testing features for real-time insights and proactive issue resolution, ensuring data availability and reliability.

Evolve Beyond Backup

Day 2 Operations

Now that your AWS backup strategy is implemented, you need to maintain and continually enhance it. When done correctly with the right tools, you can achieve time and cost savings. This provides you with extra capacity to focus on strategic initiatives. Getting this step right is key to unlocking the full value of your chosen solution so you can stay up to date on the latest features, better secure your environment against evolving cyberthreats, and automate more so you can focus on other strategic initiatives.

Day 2 Operations commonly include:



Maintain

Maintenance of a backup tool and environment is rarely a consideration when selecting a solution, but is critical for success in production, especially in the ever-changing cloud ecosystem. Patches, upgrades, enhancements, and more are essential components of software maintenance to take advantage of the latest and greatest while also addressing vulnerabilities.



Monitor

Staying on top of your AWS backup environment hinges on comprehensive monitoring, reporting, and alerting. Common practices include identifying protected/unprotected workloads and data, backup completion success and failure metrics, storage consumption, and more. Less common — but equally important — practices include proactive performance measurements and suggestions that aid in optimization.



Optimize

To maximize the benefits of your AWS backup solution, you need to continually optimize it. This can range from the more obvious improvements, like balancing performance and costs to improve RPOs and RTOs while freeing up budgets, to utilizing APIs present in the solution to further automate processes and/or integrate with other applications and software.

Plan for the Future

Unless your organization was born-in-the-cloud and standardized on AWS, it's incredibly likely that you had or still have an on-premises footprint. Of equal likelihood is that some of those on-premises workloads were lifted and shifted to the cloud, or rearchitected for cloud-native, moving from a physical or virtualized platform to AWS. The idea here to acknowledge the diverse nature of your infrastructure (e.g., hybrid or multi cloud), and emphasize the importance of being able to move your workloads between platforms without being tied to one solution. While AWS may or may not be the ultimate destination for your workloads today, the ability to have control and the option to move your workloads elsewhere remains relevant.

Some examples where workload portability is essential include:



Workload transformation

Migrating a workload from IaaS to PaaS, for example an SQL database running on Amazon EC2 to Amazon RDS.



Disaster recovery

Temporarily recovering workloads to another AWS region or another platform in the event of a regional outage.



Dev/test

Utilizing production fresh data stored in backups of workloads in one account, region, or platform to another.



Avoiding platform lock-in

Migrating workloads and data from one platform to another, e.g., on-premises to cloud, cloud to cloud, cloud to on-premises.



Veeam offers centralized management and monitoring for hybrid environments, ensuring efficient data protection and management. It provides valuable insights and recommendations based on best practices, helping organizations optimize resources, forecast storage needs, and identify cost-saving opportunities. With comprehensive reporting and analytics, Veeam enables organizations to plan strategically for future growth and scalability in the cloud.

Summary

In this ever-changing cloud landscape, it can be challenging to determine the right backup and recovery strategy for your organization.

Veeam delivers the #1 Global Data Resilience solution that covers not only your data on AWS, but your entire hybrid and multi-cloud environment. With native, secure, policy-based protection, you can have confidence in reliable recovery from accidental deletion, ransomware, and other data loss scenarios. An API-first approach, immutable backups and full- and file-level restores ensure resilient protection that's easy and cost-optimized, freeing up time and resources for strategic IT priorities.

What's more, Veeam's AWS-native solution seamlessly integrates with other cloud, virtual, physical, SaaS and Kubernetes workloads under a single platform, standardizing backup, recovery, and monitoring across the hybrid- and multi cloud. Plus, with Data Freedom, you have complete control over your data, allowing you to move it easily between environments without being locked into one platform.

About Veeam Software

Veeam®, the #1 global market leader in data resilience, believes every business should be able to bounce forward after a disruption with the confidence and control of all their data whenever and wherever they need it. Veeam calls this radical resilience, and we're obsessed with creating innovative ways to help our customers achieve it. Veeam solutions are purpose-built for powering data resilience by providing data backup, data recovery, data freedom, data security, and data intelligence. With Veeam, IT and security leaders rest easy knowing that their apps and data are protected and always available across their cloud, virtual, physical, SaaS, and Kubernetes environments. Headquartered in Seattle with offices in more than 30 countries, Veeam protects over 550,000 customers worldwide, including 74% of the Global 2000, that trust Veeam to keep their businesses running. Radical resilience starts with Veeam. Learn more at www.veeam.com or follow Veeam on LinkedIn [@veeam-software](https://www.linkedin.com/company/veeam) and X [@veeam](https://twitter.com/veeam).

→ To see Veeam with AWS in action [click here](#) to watch our demo series.