AWS Cost Optimization and Best Practices Guide

This whitepaper offers AWS Cost Optimization and Best Practices for organizing, managing, securing, and optimizing your AWS infrastructure in the Cloud.

These best practices are focused on organizations that purchase their Cloud Platform subscription through AWS and not an AWS Solution Provider Partner. This is because the organization has much more control of pricing when dealing directly with AWS.

Direct vs. Indirect Purchases

When purchasing through an AWS Solution Provider Partner, the organization is only dealing with the Partner for pricing and not AWS. Organizations have much less control of negotiating pricing with Partners who are providing them with an AWS account through their organization.

Third Party Licensing Not Included

This paper also only addresses AWS costs and not the cost of any third-party licenses that an organization is running within the AWS Cloud. If you are currently utilizing vendor licenses on-premise, do not assume that those licenses are transferrable to an AWS Cloud environment for the same use and allocation of licensing.

Ensuring adherence to the following AWS Cost Optimization & Best Practices goes a long way in successfully transitioning to AWS, and maintaining an optimized and secure environment once you get there:

1. Optimize your Usage and Costs
2. Understand AWS Support Plans
3. Organize your AWS Environment
4. Adopt the AWS Well-Architected Framework
5. Tag your AWS Resources
6. Secure your AWS Environment
7. Go Serverless and Automate

Note: *There are important AWS services and tools that most organizations don't fully utilize, thus marring the security and efficiency of their AWS Cloud operations. Some of these services and tools are free, and others cost something, with the cost outweighed by the immense benefits they bring to the table. Adopt as many of these tools as possible.*

Table Of Contents

Optimize your Usage and Costs

A move to the Cloud may not be necessarily cheaper than running your IT workloads on-premise. With that in mind, it's important to identify the ways you can save money when you move workloads to a Cloud platform.

You can optimize your AWS usage and costs in two ways: reduce your usage by right-sizing instances and turning off instances you aren't using, and by purchasing instances at discounted prices. Reduce your bills by taking advantage of the generous discounting offered by AWS on its EC2 instance, RDS databases, and other services.

Often organizations spend more than they should by using EC2 instances and EBS volumes that are larger than they need. Right-sizing is how you correctly size those instances and volumes so that you pay less for the same workloads. It's common to find that you're running your instances, especially in development and staging environments, during holidays or after working hours. You need to find all those instances and turn them off when they aren't in use. If you have EBS volumes that aren't in use, make a backup (snapshot) of the volumes and remove those volumes to save money.

AWS Discounts

There are two major types of AWS discounts – private and public. Private discounts are specific agreements you negotiate with AWS, with the discount levels depending on how much you spend on AWS services. Public discounts are available to all AWS users, without explicit agreements or contracts with AWS. These are well-known discounts such as Savings Plans and Reserved Instances.

Private Discounts

For situations where an organization is prepared to commit to a certain level of spend for Amazon Web Services, you may negotiate with AWS for a discount based on your spending levels. You can negotiate a Private Pricing Addendum (PPA) or Private Pricing Term Agreement with AWS.

A PPA has two components:

- The most common component is a discount program (previously called Enterprise Discount Program (EDP)) that applies across all AWS services and is tied to your annual spending commitment. The discount rate is negotiable between each customer and AWS. There is also a well-known credits program called the Migration Acceleration Program (MAP) to incentivize your move to the AWS Cloud. Normally, AWS looks for a minimum of half a million annual spend before offering a PPA.

- For very large customers, AWS also offers special pricing or discounts on specific AWS services for large-volume commitments. For example, if you're using petabytes of Amazon S3 storage and envisage that usage growth in the future, you can negotiate special lower prices for the storage, at a large discount from the published S3 prices.

A PPA is NOT required to be in place in order to use AWS. However, the size of the discount relies heavily on the total amount of spend. Spending more than $500k and less than $1 million may get an organization an additional 6%-7%.

Public Discounts

Public discounting has options that an AWS customer can choose to implement without any formal authorization. The discounts are the same for everyone, but they have requirements that need to be met. The following options are all based around Amazon EC2 Instances and can be purchased in 1 or 3-year terms unless noted otherwise:

- **Reserved Instances** can save up to 72% over standard on-demand prices. They support Amazon EC2, RDS (Relational Database Service), Amazon ElastiCache, Amazon DynamoDB, and Amazon Redshift. They require a committed utilization of the reserved EC2 instance, RDS database, and so on.

- **EC2 Instance Savings Plans** can save up to 72% over standard on-demand prices. They are only available for use with a specific instance family within the same region. You can switch between Windows and Linux within the covered region. They require a committed spend amount.

- **Compute Savings Plans** can save up to 66% over standard on-demand prices. They support EC2 instances regardless of region, instance family, operating system, or tenancy, including those that are part of EMR, ECS or EKS cluster. They require a committed spend amount.

- **Spot Instances** can save up to 90% over standard on-demand prices. There is no term-based commitment required. These instances can be terminated by AWS with very short notice, so they are only meant for particular uses that can handle interruptions in service.

It is important to note that should you purchase Reserved Instances or any Savings Plans and create new instances that may fall outside of the covered configurations, then the pricing will revert to on-demand for that particular instance or service.

The Right Strategy to Lower Costs

The PPA option offers discounts, but the discount rates are generally much smaller than what you get through public discounts like Savings Plans. However, private discounting covers all your spending on AWS services, and not just EC2 instances or RDS databases which are covered by public discounts.

A good strategy to reduce costs is as follows:

1. Optimize usage through right-sizing, adoption of latest generation resources, turning off instances that aren't in use, and so on.

2. Take advantage of public discounts such as Savings Plans and Reserved Instances.

3. If you're a fairly large spender (half a million dollars or more per year), negotiate a private agreement with AWS to get a discount on top of the savings you've already realized from the first two steps.

AWS PPAs offer discounts ranging from 7% or less for a $1 M commitment, to 18-20% when you commit to a spend of $20 M or more per year. That is a large amount of spend for many organizations, especially if you are newly shifting some solutions to a Cloud Platform. Unfortunately, it has been our experience that many organizations tend to stop looking for additional discounting past this stage. In the short term, it may make sense to hold off on pursuing more discounting until an organization gets a better handle on its usage requirements but should not go beyond a year or two at most.

It is also important to note that private agreements are only available when you have a direct relationship with AWS. If you're purchasing AWS services through an AWS Solution Provider partner or a Managed Service Provider, you work out an agreement with that provider directly. You are likely choosing an IT solution provided by a partner that happens to run within AWS or you are choosing a vendor that controls all of your connectivity with AWS.

When you work through a third-party AWS service provider, you do not have a contractual relationship with AWS, and so, you would not be able to negotiate any public or private discount agreements with AWS. You would have to discuss any overall pricing for the service with the provider.

When you are starting out with an AWS environment, you have an opportunity to negotiate a Private Pricing Agreement (PPA). However, environments that spend less than $500k per year do not enable discounting, and for those that spend between $500k to $1 million, the discounting is very limited. This is why it is so critical to first pursue public discounting through AWS EC2 Compute and EC2 Savings Plans, RDS Reserved Instances, and similar discounting strategies.

Note: *AWS estimates that you can reduce your bills by up to 70% by shutting down non-production instances during off hours.*

AWS Tools for Cost and Usage Optimization

Fortunately, you don't necessarily need to use third-party solutions to gain insights into your AWS usage and learn how to optimize your usage. AWS offers several powerful cost and optimization tools, the most important of which are listed here:

- **AWS Cost Explorer:** Helps you visualize and manage your AWS costs from the AWS Console. You can generate reports that help you understand what is driving your AWS costs and usage. You can also forecast your costs and usage over time.
- **AWS Trusted Advisor:** Helps reduce your costs by checking to see if you're following AWS best practice recommendations and offers recommendations to optimize your AWS infrastructure.

- **AWS Cost and Usage Reports (CUR):** CUR is the most comprehensive set of AWS usage and cost data. You can mine the CUR data using custom SQL statements to generate insights into your usage and costs. It's easy to setup CUR and run reports against it.
- **AWS Budgets:** You can use AWS budgets to get alerts when your costs hit a target spend level, and set reservation utilization and coverage targets to ensure that you're hitting those targets.

Table 1 shows how Miro's Cost Optimization Service stacks up against AWS's cost and usage-related services

FEATURES	Miro Analyzed Cost & Usage Report	AWS Budgets	AWS Cost Explorer	AWS Trusted Advisor	AWS Cost Management	AWS Compute Optimizer
Set a Budget	✓	✓	✓	✓	✓	✓
Cost Breakdown by Service	✓		✓	✓	✓	✓
Identify Unused & Underused Assets	✓			✓		
Recommend & Review Savings Plan	✓				✓	
Upsizing & Downsizing Suggestions	✓					✓
Complete View of Spending & Usage	✓					
Team of Experts who Understand your Business	✓					

Table 1. *How Miro's Cost Optimization Service compares to AWS Cost and Usage Tools*

AWS Best Practices for Optimizing Usage and Costs

To optimize your AWS usage, you must, at the minimum, follow these AWS cost and usage best practices:

- **Use correct pricing plans:** Don't buy EC2 instances, RDS database instances and ElastiCache instances (and some other resources) at sticker price! Use AWS Compute and AWS EC2 Savings Plans, and Amazon RDS and ElastiCache Reserved Instances to drastically lower your bills.
- **Use newer technologies:** Watch for infrastructure updates from AWS, as AWS frequently updates its technology offering more performance at a lower cost. Look out for opportunities to replace "previous generation" technology with "current generation" technology, a strategy referred to as "modernizing".
- **Configure storage correctly:** Ensure that you remove/release storage drives from instances after terminating the instances. Choose appropriate Amazon S3 storage classes based on your use cases, to reduce costs. AWS estimates that you can save about 50% of your storage costs by correctly choosing the data type and storage class.
- **Don't over-provision resources:** Consolidate your computing jobs to fewer EC2 instances and right-size the instances to match your workload requirements. AWS estimates that its customers saved roughly 36% of their cost by right-sizing their infrastructure.
- **Eliminate orphaned resources:** Eliminate all resources that you are getting billed for, but aren't using. The orphaned resources could be network devices or storage drives that are no longer in use.

Read your AWS Bills!

The warning to read your AWS bills isn't made lightly. It's amazing how much you can learn about your Cloud usage and costs by carefully going through your monthly AWS Invoices and the up-to-date bills provided in the AWS Console Billing page. Perusing your monthly bills carefully, line by line, helps you figure out your biggest cost drivers and also quickly zoom in on spending anomalies month-over-month.

Understand AWS Support Plans

AWS Support is a critical component of your AWS environment. Support can help you fix your infrastructure and AWS Service-related issues. Support also can help you follow AWS best practices, and help optimize your usage and costs.

AWS offers multiple Support Plans to help manage your AWS environment:

- **Developer Support:** The basic support level meant only for testing a new test AWS system. AWS charges the greater of $29/month or 3% of your monthly AWS usage. You can only access AWS during business hours with this Support level.
- **Business Support:** Offers 24/7 access to AWS Cloud Support Engineers and is the minimum recommended support tier if you're running production workloads. You pay the greater of $100/month or a percentage of your monthly usage, such as 10% of monthly AWS usage for the first $0-10K, and 5% of usage from $80K-250K.
- **Enterprise-On-Ramp Support:** Designed for business-critical systems, with a support response time of fewer than 30 minutes when such a system goes down. AWS assigns a pool of Technical Account Managers to provide proactive guidance. AWS bills you the greater of $5,500 or 10% of your monthly AWS usage cost.
- **Enterprise Support:** Another business-critical system support level, with a less than 15 minute response time when a mission-critical system goes down. AWS assigns a Designated Technical Account Manager to your account, who proactively monitors your environment and helps optimize your usage. AWS bills you $15,000 or 7% of usage if you spend $150K-500K per month or 3% of your monthly AWS usage over $1M.

If you're a small-sized organization running production workloads, get Business Support. If your workload includes mission-critical work, get either Enterprise-On-Ramp or Enterprise Support, depending on how critical the workloads are, and how much support you need from AWS in running those workloads.

Organizing an AWS Environment

When you start out with AWS in the Cloud, you usually start small, and may run everything from a single account, such as what users do with their personal (not business) AWS accounts. However, as your AWS usage ramps up, the single account strategy won't work any longer.

Multi-Account Strategy

You should organize your AWS environment carefully, by adopting what's known as a multi-account strategy. An AWS account could be the same as the account through which you pay your AWS invoices and bills. However, it's a best practice to create separate accounts. Similarly, you can separate for each department or division in an organization.

An AWS account holds all the AWS resources and data for an individual unit in your company. For example, your analytic platform and your Ecommerce website could be organized under separate accounts, so you can separate your HR and Finance departments into their own AWS account.

Using multiple AWS accounts to manage your infrastructure enables you to demarcate various applications and departments. Organizing your environment under multiple accounts offers better security, reliability and helps you optimize your AWS costs.

How many AWS Accounts?

There's no magic number of AWS accounts that you must create in your organization, with companies having anywhere from a handful to thousands of accounts. You don't incur any cost for creating more accounts. Costs are simply based on the amount of resources you use overall across your organization, regardless of which AWS account uses those resources.

AWS Organizations

Although not mandatory, it's a best practice to create an AWS Organization at the outset. This allows you to organize all your AWS accounts under this Organization. You can then create organization units (OUs) under the root of the Organization and place the individual accounts into separate OUs.

An AWS Organization has a management account through which you create OUs, individual accounts, and policies to manage the accounts. Figure 1-1 shows a typical AWS Organization with OUs and individual AWS accounts, called member accounts. The member accounts are also called linked accounts, and the management account, the payer account. AWS recommends that you create two *Foundational OUs* – Security and Infrastructure – in your AWS Organization. In addition, you create additional OUs to hold all your individual units that use the AWS Cloud.

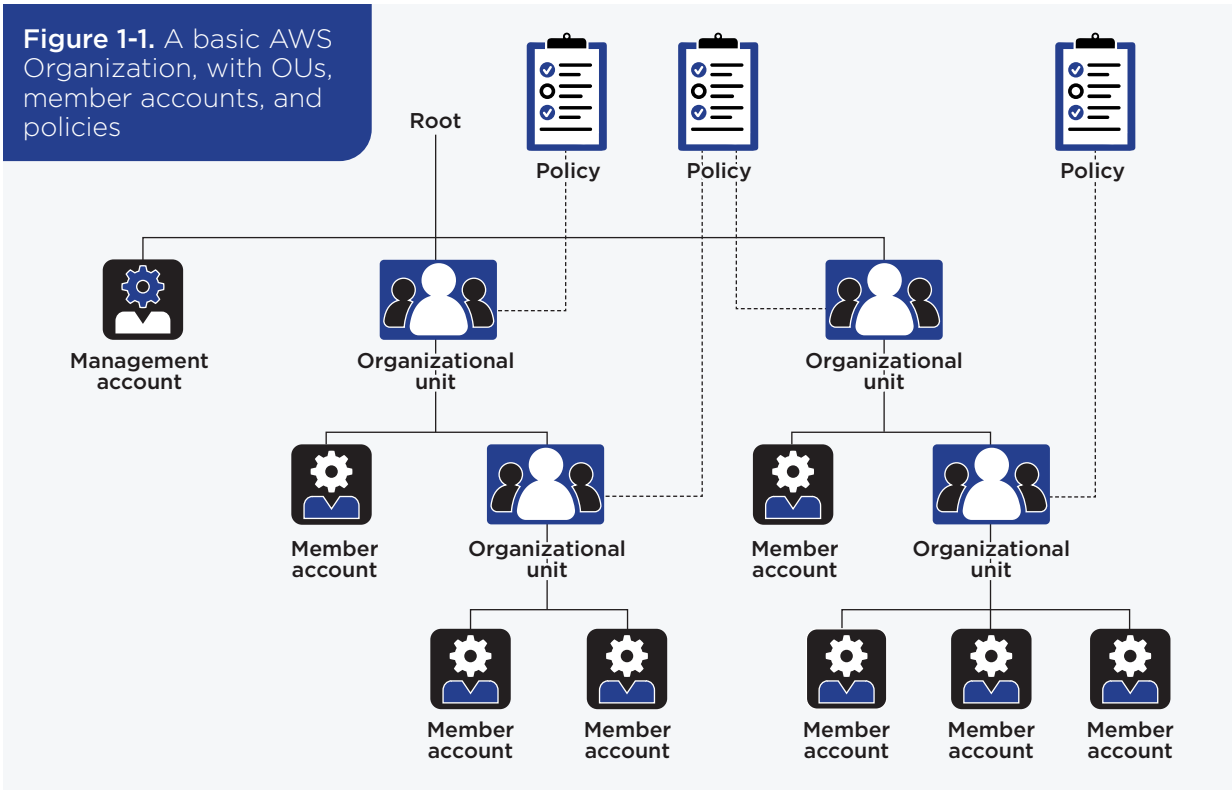

Figure 1-1. A basic AWS Organization, with OUs, member accounts, and policies

Benefits of Using a Multi-Account Strategy

Multiple AWS accounts in an organization offer the following benefits:

- **Separate your workloads:** You can improve decision making when you separate workloads into distinct accounts. This not only provides more control to individual units over their AWS usage, but also helps an organization implement separate guidance rules for security and compliance to the units.
- **Apply different security mechanisms to groups:** You can apply different control policies and security mechanisms to different workloads. For example, you can run your production workloads much more securely with stringent security and compliance polices, as compared to how you secure your development and staging environments.
- **Isolate resources:** Since the AWS resources in different AWS accounts are provisioned and managed separately from the resources in other accounts, any security breaches or issues caused by misconfiguration are limited to just a single account, leaving other accounts untouched.
- **Apply common policies:** OUs help you manage your AWS accounts by applying policies at the OU level, instead of doing so at the account level. For example, if you have 25 AWS accounts under Organizational Unit A, you can simply apply a policy at the OU level, which results in that policy being applied to all 25 of the AWS accounts in that OU.
- **Facilitate the optimization and reporting of costs:** It's easier to budget, report, and forecast AWS costs when you organize along multiple AWS accounts.

Adopt the AWS Well-Architected Framework

The AWS Well-Architected Framework (AWF) should be upper most in your mind when designing and managing the AWS Cloud.

The AWS Well-Architected Framework is a set of six principles (or pillars as AWS calls them) that ensure you're architecting, and operating a reliable, cost-effective, secure and efficient system. AWS Well-Architected Framework is free of charge, so there's no excuse for not using it.

You should review all your workloads against the best practices incorporated in the Well-Architected Framework, thereby quickly getting an idea as to how close or how far you are from these best practices. AWS offers the *AWS Well-Architected Tool,* which helps you run the reviews from the AWS Management Console. Merely answer a few simple questions, and find out how your systems are aligned with AWS best practices, and receive guidance to improve your architecture to align yourself closely with the best practices.

Note: *The AWS Well-Architected Tool now has the ability to create Custom Lenses, which help you tailor reviews to your needs, and allow you to create your own pillars, questions and best practices. You can also specify custom rules to flag issues as high or medium.*

The Six Pillars of the AWS Well-Architected Framework

Operational Excellence Pillar

The key goal here is to efficiently manage and monitor your systems in the Cloud and focus on a continual improvement of all processes. Define daily operational standards and codify your event response procedures. Automate your operations with Playbooks and Runbooks.

By standardizing operations, you can reduce risk and increase efficiency, whether you're dealing with daily operations, inventory and patch management, or responses to unexpected events.

Performance Efficiency Pillar

Focus on structuring and streamlining your resource allocation in the AWS cloud. Determine how you select the resource types and resource sizes for various workloads, with the goal of optimizing your resource usage. Right-sizing of all your Compute and other resources and terminating "zombie" resources (such as idle EC2 instances and RDS databases) falls into this category.

CloudWatch is a powerful monitoring and logging tool, and creating efficient CloudWatch Dashboards is critical to monitoring your EC2 instances.

Cost Optimization Pillar

Cost optimization is all about eliminating or reducing unnecessary costs and scaling your environment without overspending. The goal is to understand where you're spending your money and control how you allocate that money. Often, AWS provides multiple ways to perform common tasks, such as messaging, notifications and automatic remediation of issues. You must learn how to use cost-effective resources.

Part of optimizing costs is learning how to visualize the costs. Use Cloud Intelligence Dashboards and figure out how to determine workload efficiency. Working with the various AWS Console tools such as Compute Optimizer, you can find EC2 instances that are larger than you need for your workloads. Setting up the AWS Cost and Usage Repository (CUR), the most comprehensive source for your AWS costs and usage, is part of the Cost Optimization pillar.

Reliability Pillar

Organizations must have procedures and plans in place to quickly recover from intermittent failures as well as major disasters that threaten to put their systems out of commission. The Reliability Pillar encompasses the following related components of a smoothly running system:

- **Reliability:** How a workload performs correctly and consistently
- **Resiliency:** How a workload recovers quickly from disruptions such as a network issue or a misconfigured application code

To ensure resiliency, you must replicate data across regions, and continuously test the resilience of your EC2 instances as well as your backup and restore plans for critical data. You also implement health checks throughout your system to proactively manage interruptions. Use AWS's Chaos Engineering tools to test the resilience of your servers and databases. Test application resilience with the AWS Resiliency Hub.

You must also set up disaster recovery plans and take advantage of the AWS Elastic Disaster Recovery service to automate recovery from disasters. Depending on how critical the systems are, you may opt for basic disaster recovery by restoring your backups or set up a warm or hot standby system to quickly takeover a downed critical workload.

Sustainability Pillar

The Sustainability Pillar is geared to how you minimize the environmental impact of being in the AWS Cloud. The key idea here is that you and AWS share the responsibility of sustainability in the Cloud. Figure 1-2 shows how you, the customer, and AWS share the responsibility.

Figure 1-2. How you and AWS share the Cloud Sustainability responsibilites

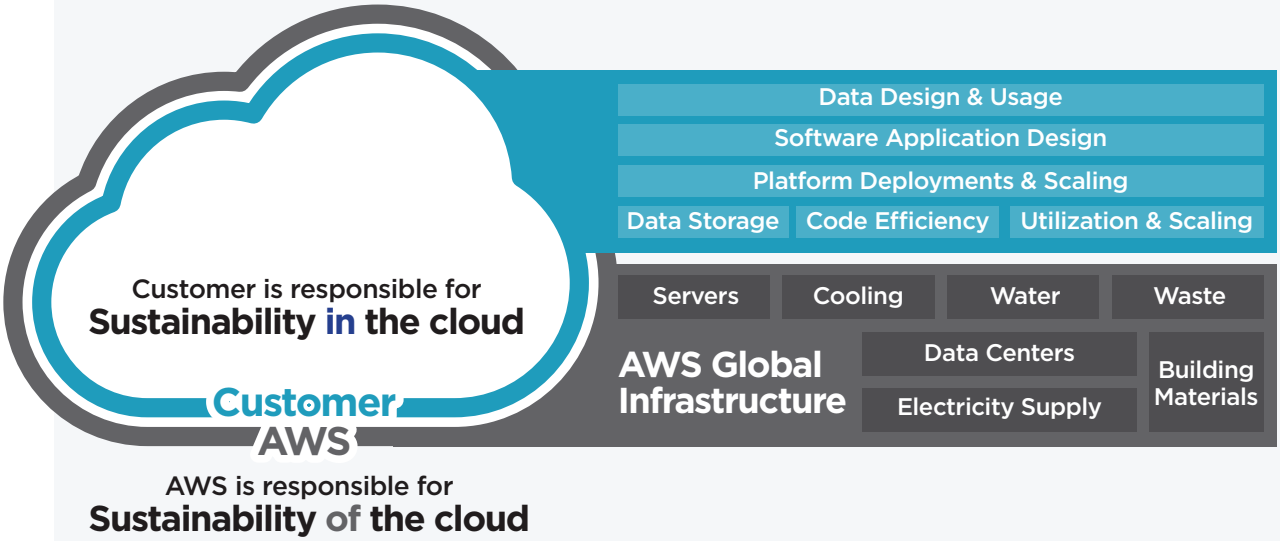

As Figure 1-2 shows, AWS takes care of the infrastructure in the data centers, including power, cooling and water. You are responsible for efficiently designing and coding your applications, proper utilization of the infrastructure and scaling efficiently.

Running your workloads on new and efficient hardware (current generation instead of previous generation servers, for example), including, if necessary, the migration of workloads to the latest generation servers such as Graviton, fall under this pillar.

Security Pillar

As with the Sustainability Pillar, you and AWS share the responsibility to support the Security Pillar. AWS secures the Cloud infrastructure and data centers, and you implement security for your databases, data sets, user credentials, and permissions. You also must set up your own SIEM (Security Information and Event Management) or other systems to detect and act on security events.

Please see the major section *"Secure your AWS Environment"* to learn more about securing your systems in the AWS Cloud.

Note: *AWS offers free labs that help you learn how to implement all six pillars of the AWS Well-Architected Framework.*

Here's the link: https://www.wellarchitectedlabs.com/.

Tag your AWS Resources

In an AWS account, there are many linked accounts, and many groups under these accounts. AWS doesn't know which user from which of the groups has created the AWS resources that you use. For example, if user A from Linked Account 12345678 creates an AWS instance, that instance has no information attached to it that marks it as the instance created by User A from Account 12345678. So, how do you track these instances and other AWS resources? You use tags for labeling all your resources. Tags come free of charge, and are easy to apply to your resources, either during their creation, or later.

What's an AWS Tag?

A tag for a resource is how you associate a set of labels that you create (metadata) to the resources, so you can easily identify the user, department, or cost center that's responsible for the creation and usage of that resource. Tagging your resources is how you maintain enterprise-level visibility and control of your AWS resources.

How you Use Tags

AWS tags serve a variety of purposes, such as controlling access to the resources, Cloud financial management, and automation of AWS-related activities. Tagging helps you group together a bunch of resources, thereby facilitating actions such as patching all EC2 instances running in the production environment, or all resources pertaining to a specific application or workload.

Tagging Policies

Create tagging standards for your environment, including a set of mandatory and discretionary tags. Common tags include items such as environment, owner, department, and cost center.

Always, always tag all taggable resources (most AWS resources are taggable)! Without tagging, you can't identify who owns and runs which resources, making it near impossible to optimize and manage the many AWS resources that you use.

A common tagging strategy is to assign default pre-defined tags to resources that aren't explicitly tagged.

Enforce Tagging

After creating a tagging strategy, you must establish controls to ensure that new resources are tagged when they're created, thereby automating the tagging process. Through automation, you can find resources with missing or non-compliant tags and take action to remediate the lapses.

Secure your AWS Cloud Systems

As mentioned earlier, you're fully responsible for securing your data and users, as well as protecting yourself from common internet security threats. AWS doesn't charge you for managing access to your AWS systems via IAM (Identity and Access Management), or for using AWS Organizations to centrally manage your Cloud environment. Amazon CloudWatch is a key tool for observing your resource usage and monitoring your applications. You can set up CloudWatch to automatically respond to system-wide performance changes.

AWS offers many paid security related services, most of which you should at the very least, use in your production environments. The following is a summary of these security services.

Threat Detection and Incident Response

Use the *AWS Security Hub* to automate security best practice checks, and even put in place automatic remediation of security issues to which you are alerted by the checks.

Amazon GuardDuty services detects malicious activity to protect your AWS accounts and data. *Amazon Inspector*, on the other hand, scans your workloads for software vulnerabilities as well as unwarranted network exposure.

Amazon Detective quickly analyzes and identifies the root causes of potential security issues or any suspicious activity in your AWS system.

Application Security

Two AWS services – *AWS WAF* (Web Application Firewall) and *AWS Shield* go a long way in protecting applications running in the AWS Cloud.

AWS WAF is a firewall that protects your web applications against common web exploits. AWS Shield (comes in two varieties – free and paid) protects your systems against the well-known *Distributed Denial of Service* (DDoS) attacks that could bring your system to a halt.

Data Protection

There are two basic data protection features that you're very likely to use. First, is *AWS Key Management Service* (KMS) which helps you create and manage cryptographic keys that you can use across several AWS services. Second is *AWS Private Certificate Authority*, and that helps you secure your applications and servers.

Two important data protection services that you must explore are *Amazon Macie* and *AWS Secrets Manager.* Amazon Macie uses machine learning to both discover and protect sensitive data that you store in your AWS environment. AWS Secrets Manager helps you safeguard and manage secrets such as database credentials.

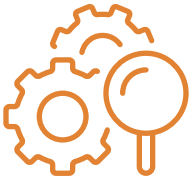

Compliance

Besides the basically free AWS Cloud Trail service that enables operational, compliance, and risk auditing in your account, there are two other services that you must investigate to strengthen your compliance posture:

- **AWS Artifact:** This free service lets you access AWS's security and compliance reports. So, these are AWS's own reports, which you can then forward to any third-party organization that's auditing your organization.
- **AWS Audit Manager:** Helps simplify how you assess risk and compliance by continuously auditing your AWS usage.

Go Serverless and Automate

Traditional computing involves server management. You can still do this in the Cloud, but it usually isn't a smart idea to set yourself up for server management. AWS offers serverless technologies based on a pay-for-value billing model that helps you manage your data, perform computing tasks and run code, without worrying about maintenance activities like patching, and capacity provisioning.

Note: *Serverless technology means you only pay for what you use - there's no infrastructure to pay for. This means that you automatically optimize your usage and there's never an over-provisioning of resources!*

Adopt AWS's serverless technologies where you can, to increase your agility, without getting enmeshed in complex infrastructure management.

Types of Serverless Technologies

Several of AWS's serverless technologies help you architect modern microservices. The following is a summary of the serverless services you may want to consider.

Computing

There are two modern and powerful compute serverless technologies:

- **AWS Lambda** is a foundational serverless compute service that enables you to run your application code without having to provision any servers whatsoever. You just pay for the brief time your code snippets run in the AWS Cloud.
- **AWS Fargate:** This is a serverless compute engine that helps you run containers and Kubernetes clusters in the Cloud.

Data Stores

You can setup and run your own databases on EC2 instances, but there are better alternatives:

- **Amazon DynamoDB:** A highly scalable and fast key-value and document database service.
- **Amazon Aurora Serverless:** You can use either MySQL or PostgreSQL-compatible Serverless relational databases that automatically scale depending on your application needs.
- **Amazon Redshift Serverless:** Allows you to run powerful, scalable analytical applications without setting up and managing a data warehouse.
- **Amazon Elastic File System (EFS):** You can set up scalable shared storage.

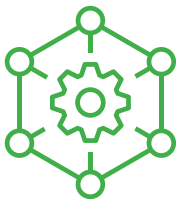

Application Integration

It's often the case that you need to set up communication between your applications. In the AWS Cloud, you don't need to integrate your applications using custom code and third-party tools. Use the following AWS services to smoothly integrate your applications.

- **Amazon EventBridge:** Helps you build event-driven applications without provisioning any infrastructure.
- **Amazon Simple Queue Service (SQS):** This is a message queue service that helps you decouple your microservices.
- **Amazon Simple Notification Service (SNS):** This managed messaging service helps applications send event-driven notifications to other applications or users in your Cloud.
- **AWS Step Functions:** Powerful workflow orchestrator that enables you to sequentially tie together AWS services in your applications.

Automate to the Max

Reduce manual operations by automating system operations wherever possible. Use *AWS Systems Manager* and *EC2 Image Builder* to automate instance patching. You can automate IAM user cleanup, deployment of VPCs, Web applications, Detective Controls, management of IAM Groups and Roles, Web Application Firewalls, and many other components of your system.

Use the AWS provided free *CloudFormation* service to model and set up your AWS resources, reducing manual, error prone work. AWS CloudFormation helps you quickly replicate your infrastructure, and easily track and control infrastructure changes.

In addition to AWS CloudFormation, use tools such as Ansible and Terraform to build and manage your infrastructure.