

Security Pillar

AWS Well-Architected Framework

July 2020



Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents current AWS product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

© 2020 Amazon Web Services, Inc. or its affiliates. All rights reserved.

Contents

- Introduction 1
- Security 2
 - Design Principles..... 2
 - Definition..... 3
 - Operating Your Workload Securely 3
 - AWS Account Management and Separation..... 5
- Identity and Access Management..... 7
 - Identity Management..... 7
 - Permissions Management 11
- Detection..... 15
 - Configure 15
 - Investigate 18
- Infrastructure Protection 19
 - Protecting Networks 20
 - Protecting Compute 23
- Data Protection 27
 - Data Classification 27
 - Protecting Data at Rest..... 29
 - Protecting Data in Transit 32
- Incident Response..... 34
 - Design Goals of Cloud Response 34
 - Educate..... 35
 - Prepare..... 36
 - Simulate..... 38
 - Iterate..... 39
- Conclusion 40

Contributors	40
Further Reading.....	41
Document Revisions.....	41

Abstract

The focus of this paper is the security pillar of the [Well-Architected Framework](#). It provides guidance to help you apply best practices, current recommendations in the design, delivery, and maintenance of secure AWS workloads.

Introduction

The [AWS Well-Architected Framework](#) helps you understand trade-offs for decisions you make while building workloads on AWS. By using the Framework, you will learn current architectural best practices for designing and operating reliable, secure, efficient, and cost-effective workloads in the cloud. It provides a way for you to consistently measure your workload against best practices and identify areas for improvement. We believe that having well-architected workloads greatly increases the likelihood of business success.

The framework is based on five pillars:

- Operational Excellence
- Security
- Reliability
- Performance Efficiency
- Cost Optimization

This paper focuses on the security pillar. This will help you meet your business and regulatory requirements by following current AWS recommendations. It's intended for those in technology roles, such as chief technology officers (CTOs), chief information security officers (CSOs/CISOs), architects, developers, and operations team members.

After reading this paper, you will understand AWS current recommendations and strategies to use when designing cloud architectures with security in mind. This paper doesn't provide implementation details or architectural patterns but does include references to appropriate resources for this information. By adopting the practices in this paper, you can build architectures that protect your data and systems, control access, and respond automatically to security events.

Security

The security pillar describes how to take advantage of cloud technologies to protect data, systems, and assets in a way that can improve your security posture. This paper provides in-depth, best-practice guidance for architecting secure workloads on AWS.

Design Principles

In the cloud, there are a number of principles that can help you strengthen your workload security:

- **Implement a strong identity foundation:** Implement the principle of least privilege and enforce separation of duties with appropriate authorization for each interaction with your AWS resources. Centralize identity management, and aim to eliminate reliance on long-term static credentials.
- **Enable traceability:** Monitor, alert, and audit actions and changes to your environment in real time. Integrate log and metric collection with systems to automatically investigate and take action.
- **Apply security at all layers:** Apply a defense in depth approach with multiple security controls. Apply to all layers (for example, edge of network, VPC, load balancing, every instance and compute service, operating system, application, and code).
- **Automate security best practices:** Automated software-based security mechanisms improve your ability to securely scale more rapidly and cost-effectively. Create secure architectures, including the implementation of controls that are defined and managed as code in version-controlled templates.
- **Protect data in transit and at rest:** Classify your data into sensitivity levels and use mechanisms, such as encryption, tokenization, and access control where appropriate.
- **Keep people away from data:** Use mechanisms and tools to reduce or eliminate the need for direct access or manual processing of data. This reduces the risk of mishandling or modification and human error when handling sensitive data.
- **Prepare for security events:** Prepare for an incident by having incident management and investigation policy and processes that align to your organizational requirements. Run incident response simulations and use tools with automation to increase your speed for detection, investigation, and recovery.

Definition

Security in the cloud is composed of five areas:

1. Identity and access management
2. Detection
3. Infrastructure protection
4. Data protection
5. Incident response

Security and Compliance is a shared responsibility between AWS and you, the customer. This shared model can help reduce your operational burden. You should carefully examine the services you choose as your responsibilities vary depending on the services used, the integration of those services into your IT environment, and applicable laws and regulations. The nature of this shared responsibility also provides the flexibility and control that permits the deployment.

Operating Your Workload Securely

To operate your workload securely, you must apply overarching best practices to every area of security. Take requirements and processes that you have defined in operational excellence at an organizational and workload level, and apply them to all areas. Staying up to date with AWS and industry recommendations and threat intelligence helps you evolve your threat model and control objectives. Automating security processes, testing, and validation allow you to scale your security operations.

Identify and prioritize risks using a threat model: Use a threat model to identify and maintain an up-to-date register of potential threats. Prioritize your threats and adapt your security controls to prevent, detect, and respond. Revisit and maintain this in the context of the evolving security landscape.

Identify and validate control objectives: Based on your compliance requirements and risks identified from your threat model, derive and validate the control objectives and controls that you need to apply to your workload. Ongoing validation of control objectives and controls help you measure the effectiveness of risk mitigation.

Keep up to date with security threats: Recognize attack vectors by staying up to date with the latest security threats to help you define and implement appropriate controls.

Keep up to date with security recommendations: Stay up to date with both AWS and industry security recommendations to evolve the security posture of your workload.

Evaluate and implement new security services and features regularly: Evaluate and implement security services and features from AWS and APN Partners that allow you to evolve the security posture of your workload.

Automate testing and validation of security controls in pipelines: Establish secure baselines and templates for security mechanisms that are tested and validated as part of your build, pipelines, and processes. Use tools and automation to test and validate all security controls continuously. For example, scan items such as machine images and infrastructure as code templates for security vulnerabilities, irregularities, and drift from an established baseline at each stage.

Reducing the number of security misconfigurations introduced into a production environment is critical—the more quality control and reduction of defects you can perform in the build process, the better. Design continuous integration and continuous deployment (CI/CD) pipelines to test for security issues whenever possible. CI/CD pipelines offer the opportunity to enhance security at each stage of build and delivery. CI/CD security tooling must also be kept updated to mitigate evolving threats.

Resources

Refer to the following resources to learn more about operating your workload securely.

Videos

- [Security Best Practices the Well-Architected Way](#)
- [Enable AWS adoption at scale with automation and governance](#)
- [AWS Security Hub: Manage Security Alerts & Automate Compliance](#)
- [Automate your security on AWS](#)

Documentation

- [Overview of Security Processes](#)
- [Security Bulletins](#)
- [Security Blog](#)
- [What's New with AWS](#)
- [AWS Security Audit Guidelines](#)



- [Set Up a CI/CD Pipeline on AWS](#)

AWS Account Management and Separation

We recommend that you organize workloads in separate accounts and group accounts based on function, compliance requirements, or a common set of controls rather than mirroring your organization's reporting structure. In AWS, accounts are a hard boundary, zero trust container for your resources. For example, account-level separation is strongly recommended for isolating production workloads from development and test workloads.

Separate workloads using accounts: Start with security and infrastructure in mind to enable your organization to set common guardrails as your workloads grow. This approach provides boundaries and controls between workloads. Account-level separation is strongly recommended for isolating production environments from development and test environments, or providing a strong logical boundary between workloads that process data of different sensitivity levels, as defined by external compliance requirements (such as PCI-DSS or HIPAA), and workloads that don't.

Secure AWS accounts: There are a number of aspects to securing your AWS accounts, including the securing of, and not using the [root user](#), and keeping the contact information up to date. You can use [AWS Organizations](#) to centrally manage and govern your accounts as you grow and scale your workloads. AWS Organizations helps you manage accounts, set controls, and configure services across your accounts.

Manage accounts centrally: AWS Organizations [automates AWS account creation and management](#), and control of those accounts after they are created. When you create an account through AWS Organizations, it is important to consider the email address you use, as this will be the root user that allows the password to be reset. Organizations allows you to group accounts into [organizational units \(OUs\)](#), which can represent different environments based on the workload's requirements and purpose.

Set controls centrally: Control what your AWS accounts can do by only allowing specific services, Regions, and service actions at the appropriate level. AWS Organizations allows you to use service control policies (SCPs) to apply permission guardrails at the organization, organizational unit, or account level, which apply to all [AWS Identity and Access Management \(IAM\)](#) users and roles. For example, you can apply an SCP that restricts users from launching resources in Regions that you have not explicitly allowed. AWS Control Tower offers a simplified way to set up and govern multiple accounts. It automates the setup of accounts in your AWS Organization,



automates provisioning, applies [guardrails](#) (which include prevention and detection), and provides you with a dashboard for visibility.

Configure services and resources centrally: AWS Organizations helps you configure [AWS services](#) that apply to all of your accounts. For example, you can configure central logging of all actions performed across your organization using [AWS CloudTrail](#), and prevent member accounts from disabling logging. You can also centrally aggregate data for rules that you've defined using [AWS Config](#), enabling you to audit your workloads for compliance and react quickly to changes. AWS CloudFormation [StackSets](#) allow you to centrally manage AWS CloudFormation stacks across accounts and OUs in your organization. This allows you to automatically provision a new account to meet your security requirements.

Resources

Refer to the following resources to learn more about AWS recommendations for deploying and managing multiple AWS accounts.

Videos

- [Managing and governing multi-account AWS environments using AWS Organizations](#)
- [AXA: Scaling adoption with a Global Landing Zone](#)
- [Using AWS Control Tower to Govern Multi-Account AWS Environments](#)

Documentation

- [Establishing your best practice AWS environment](#)
- [AWS Organizations](#)
- [AWS Control Tower](#)
- [Working with AWS CloudFormation StackSets](#)
- [How to use service control policies to set permission guardrails across accounts in your AWS Organization](#)

Hands-on

- Lab: [AWS Account and Root User](#)



Identity and Access Management

To use AWS services, you must grant your users and applications access to resources in your AWS accounts. As you run more workloads on AWS, you need robust identity management and permissions in place to ensure that the right people have access to the right resources under the right conditions. AWS offers a large selection of capabilities to help you manage your human and machine identities and their permissions. The best practices for these capabilities fall into two main areas:

- Identity management
- Permissions management

Identity Management

There are two types of identities you need to manage when approaching operating secure AWS workloads.

- **Human Identities:** The administrators, developers, operators, and consumers of your applications require an identity to access your AWS environments and applications. These can be members of your organization, or external users with whom you collaborate, and who interact with your AWS resources via a web browser, client application, mobile app, or interactive command-line tools.
- **Machine Identities:** Your workload applications, operational tools, and components require an identity to make requests to AWS services, for example, to read data. These identities include machines running in your AWS environment, such as Amazon EC2 instances or AWS Lambda functions. You can also manage machine identities for external parties who need access. Additionally, you might also have machines outside of AWS that need access to your AWS environment.

Rely on a centralized identity provider: For workforce identities, rely on an identity provider that enables you to manage identities in a centralized place. This makes it easier to manage access across multiple applications and services, because you are creating, managing, and revoking access from a single location. For example, if someone leaves your organization, you can revoke access for all applications and services (including AWS) from one location. This reduces the need for multiple credentials and provides an opportunity to integrate with existing human resources (HR) processes.

For federation with individual AWS accounts, you can use centralized identities for AWS with a [SAML 2.0](#)-based provider with AWS IAM. You can use any provider—whether hosted by you in AWS, external to AWS, or supplied by the AWS Partner Network (APN)—that is compatible with the SAML 2.0 protocol. You can use federation between your AWS account and your chosen provider to grant a user or application access to call AWS API operations by using a SAML assertion to get temporary security credentials. Web-based single sign-on is also supported, allowing users to sign in to the AWS Management Console from your sign in portal.

For federation to multiple accounts in your AWS Organization, you can configure your identity source in [AWS Single Sign-On \(AWS SSO\)](#), and specify where your users and groups are stored. Once configured, your identity provider is your source of truth, and information can be [synchronized](#) using the System for Cross-domain Identity Management (SCIM) v2.0 protocol. You can then look up users or groups and grant them single sign-on access to AWS accounts, cloud applications, or both.

AWS SSO integrates with AWS Organizations, which enables you to configure your identity provider once and then [grant access to existing and new accounts](#) managed in your organization. AWS SSO provides you with a default store, which you can use to manage your users and groups. If you choose to use the AWS SSO store, create your users and groups and assign their level of access to your AWS accounts and applications, keeping in mind the best practice of least privilege. Alternatively, you can choose to [Connect to Your External Identity Provider](#) using SAML 2.0, or [Connect to Your Microsoft AD Directory](#) using AWS Directory Service. Once configured, you can sign into the AWS Management Console, command line interface, or the AWS mobile app, by authenticating through your central identity provider.

For managing end-users or consumers of your workloads, such as a mobile app, you can use [Amazon Cognito](#). It provides authentication, authorization, and user management for your web and mobile apps. Your users can sign in directly with a user name and password, or through a third party, such as Amazon, Apple, Facebook, or Google.

Leverage user groups and attributes: As the number of users you manage grows, you will need to determine ways to organize them so that you can manage them at scale. Place users with common security requirements in groups defined by your identity provider, and put mechanisms in place to ensure that user attributes that may be used for access control (for example, department or location) are correct and updated. Use these groups and attributes to control access, rather than individual users. This allows you to manage access centrally by changing a user's group membership or

attributes once with a [permission set](#), rather than updating many individual policies when a user's access needs change. You can use AWS SSO to manage user groups and attributes. AWS SSO supports most commonly used attributes whether they are entered manually during user creation or automatically provisioned using a synchronization engine, such as defined in the System for Cross-Domain Identity Management (SCIM) specification.

Use strong sign-in mechanisms: Enforce minimum password length, and educate your users to avoid common or reused passwords. Enforce multi-factor authentication (MFA) with software or hardware mechanisms to provide an additional layer of verification. For example, when using [AWS SSO as the identity source](#), configure the “context-aware” or “always-on” setting for MFA, and allow users to enroll their own MFA devices to accelerate adoption. When using an external identity provider (IdP), configure your IdP for MFA.

Use temporary credentials: Require identities to dynamically acquire [temporary credentials](#). For workforce identities, use AWS SSO, or federation with IAM, to access AWS accounts. For machine identities, such as EC2 instances or Lambda functions, require the use of IAM roles instead of IAM users with long term access keys.

For human identities using the AWS Management Console, require users to acquire temporary credentials and federate into AWS. You can do this using the AWS SSO user portal or configuring federation with IAM. For users requiring CLI access, ensure that they use [AWS CLI v2, which supports direct integration with AWS Single Sign-On \(AWS SSO\)](#). Users can create CLI profiles that are linked to AWS SSO accounts and roles. The CLI automatically retrieves AWS credentials from AWS SSO and refreshes them on your behalf. This eliminates the need to copy and paste temporary AWS credentials from the AWS SSO console. For SDK, users should rely on AWS STS to assume roles to receive temporary credentials. In certain cases, temporary credentials might not be practical. You should be aware of the risks of storing access keys, rotate these often, and require MFA as a condition when possible.

For cases where you need to grant consumers access to your AWS resources, use [Amazon Cognito](#) identity pools and assign them a set of temporary, limited privilege credentials to access your AWS resources. The permissions for each user are controlled through [IAM roles](#) that you create. You can define rules to choose the role for each user based on claims in the user's ID token. You can define a default role for authenticated users. You can also define a separate IAM role with limited permissions for guest users who are not authenticated.

For machine identities, you should rely on IAM roles to grant access to AWS. For EC2 instances, you can use [roles for Amazon EC2](#). You can attach an IAM role to your EC2 instance to enable your applications running on Amazon EC2 to use temporary security credentials that AWS creates, distributes, and rotates automatically. For accessing EC2 instances using keys or passwords, [AWS Systems Manager](#) is a more secure way to access and manage your instances using a pre-installed agent without the stored secret. Additionally, other AWS services, such as AWS Lambda, enable you to configure an IAM service role to grant the service permissions to perform AWS actions using temporary credentials.

Audit and rotate credentials periodically: Periodic validation, preferably through an automated tool, is necessary to verify that the correct controls are enforced. For human identities, you should require users to change their passwords periodically and retire access keys in favor of temporary credentials. We also recommend that you continuously monitor MFA settings in your identity provider. You can set up [AWS Config Rules](#) to monitor these settings. For machine identities, you should rely on temporary credentials using IAM roles. For situations where this is not possible, frequent auditing and rotating access keys is necessary.

Store and use secrets securely: For credentials that are not IAM-related, such as database logins, use a service that is designed to handle management of secrets, such as [AWS Secrets Manager](#). AWS Secrets Manager makes it easy to manage, rotate, and securely store encrypted secrets using [supported services](#). Calls to access the secrets are logged in CloudTrail for auditing purposes, and IAM permissions can grant least-privilege access to them.

Resources

Refer to the following resources to learn more about AWS best practices for protecting your AWS credentials.

Videos

- [Mastering identity at every layer of the cake](#)
- [Managing user permissions at scale with AWS SSO](#)
- [Best Practices for Managing, Retrieving, & Rotating Secrets at Scale](#)

Documentation

- [The AWS Account Root User](#)



- [AWS Account Root User Credentials vs. IAM User Credentials](#)
- [IAM Best Practices](#)
- [Setting an Account Password Policy for IAM Users](#)
- [Getting Started with AWS Secrets Manager](#)
- [Using Instance Profiles](#)
- [Temporary Security Credentials](#)
- [Identity Providers and Federation](#)

Permissions Management

Manage permissions to control access to people and machine identities that require access to AWS and your workloads. Permissions control who can access what, and under what conditions. Set permissions to specific human and machine identities to grant access to specific service actions on specific resources. Additionally, specify conditions that must be true for access to be granted. For example, you can allow developers to create new Lambda functions, but only in a specific Region. When managing your AWS environments at scale, adhere to the following best practices to ensure that identities only have the access they need and nothing more.

Define permission guardrails for your organization: As you grow and manage additional workloads in AWS, you should separate these workloads using accounts and manage those accounts using AWS Organizations. We recommend that you establish common permission guardrails that restrict access to all identities in your organization. For example, you can restrict access to specific AWS Regions, or prevent your team from deleting common resources, such as an IAM role used by your central security team. You can get started by implementing [example service control policies](#), such as preventing users from disabling key services.

You can use AWS Organizations to group accounts and set common controls on each group of accounts. To set these common controls, you can use services integrated with AWS Organizations. Specifically, you can use [service control policies \(SCPs\) to restrict access to group of accounts](#). SCPs use the IAM policy language and enable you to establish controls that all IAM principals (users and roles) adhere to. You can restrict access to specific service actions, resources and based on specific condition to meet the access control needs of your organization. If necessary, you can define exceptions

to your guardrails. For example, you can restrict service actions for all IAM entities in the account except for a specific administrator role.

Grant least privilege access: Establishing a principle of [least privilege](#) ensures that identities are only permitted to perform the most minimal set of functions necessary to fulfill a specific task, while balancing usability and efficiency. Operating on this principle limits unintended access and helps ensure that you can audit who has access to which resources. In AWS, identities have no permissions by default with the exception of the root user, which should only be used for a few [specific tasks](#).

You use policies to explicitly grant permissions attached to IAM or resource entities, such as an IAM role used by federated identities or machines, or resources (for example, S3 buckets). When you create and attach a policy, you can specify the service actions, resources, and conditions that must be true for AWS to allow access. AWS supports a variety of conditions to help you scope down access. For example, using the `PrincipalOrgID` [condition key](#), the identifier of the AWS Organizations is verified so access can be granted within your AWS Organization. You can also control requests that AWS services make on your behalf, like AWS CloudFormation creating an AWS Lambda function, by using the `CalledVia` condition key. This enables you to set granular permissions for your human and machine identities across AWS.

AWS also has capabilities that enable you to scale your permissions management and adhere to least privilege.

Permissions Boundaries: You can use permission boundaries to set the maximum permissions that an administrator can set. This enables you to delegate the ability to create and manage permissions to developers, such as the creation of an IAM role, but limit the permissions they can grant so that they cannot escalate their privilege using what they have created.

Attribute-based access control (ABAC): AWS enables you to grant permissions based on attributes. In AWS, these are called tags. Tags can be attached to IAM principals (users or roles) and to AWS resources. Using IAM policies, administrators can create a reusable policy that applies permissions based on the attributes of the IAM principal. For example, as an administrator you can use a single IAM policy that grants developers in your organization access to AWS resources that match the developers' project tags. As the team of developers adds resources to projects, permissions are automatically applied based on attributes. As a result, no policy update is required for each new resource.

Analyze public and cross account access: In AWS, you can grant access to resources in another account. You grant direct cross-account access using policies attached to resources (for example, S3 bucket policies) or by allowing an identity to assume an IAM role in another account. When using resource policies, you want to ensure you grant access to identities in your organization and are intentional about when you make a resource public. Making a resource public should be used sparingly, as this action allows anyone to access the resource. [IAM Access Analyzer](#) uses mathematical methods (that is, [provable security](#)) to identify all access paths to a resource from outside of its account. It reviews resource policies continuously, and reports findings of public and cross-account access to make it easy for you to analyze potentially broad access.

Share resources securely: As you manage workloads using separate accounts, there will be cases where you need to share resources between those accounts. We recommend that you share resources using [AWS Resource Access Manager \(AWS RAM\)](#). This service enables you to easily and securely share AWS resources within your AWS Organization and Organizational Units. Using AWS RAM, access to shared resources is automatically granted or revoked as accounts are moved in and out of the Organization or Organization Unit with which they are shared. This helps you ensure that resources are only shared with the accounts that you intend.

Reduce permissions continuously: Sometimes, when teams and projects are just getting started, you might choose to grant broad access to inspire innovation and agility. We recommend that you evaluate access continuously and restrict access to only the permissions required and achieve least privilege. AWS provides access analysis capabilities to help you identify unused access. To help you identify unused users and roles, AWS analyzes access activity and provides access key and role last used information. You can use the [last accessed timestamp](#) to [identify unused users and roles](#), and remove them. Moreover, you can review service and action last accessed information to identify and [tighten permissions for specific users and roles](#). For example, you can use last accessed information to identify the specific S3 actions that your application role requires and restrict access to only those. These feature are available in the console and programmatically to enable you to incorporate them into your infrastructure workflows and automated tools.

Establish emergency access process: You should have a process that allows emergency access to your workload, in particular your AWS accounts, in the unlikely event of an automated process or pipeline issue. This process could include a combination of different capabilities, for example, an emergency AWS cross-account

role for access, or a specific process for administrators to follow to validate and approve an emergency request.

Resources

Refer to the following resources to learn more about current AWS best practices for fine-grained authorization.

Videos

- [Become an IAM Policy Master in 60 Minutes or Less](#)
- [Separation of Duties, Least Privilege, Delegation, & CI/CD](#)

Documentation

- [Grant least privilege](#)
- [Working with Policies](#)
- [Delegating Permissions to Administer IAM Users, Groups, and Credentials](#)
- [IAM Access Analyzer](#)
- [Remove unnecessary credentials](#)
- [Assuming a role in the CLI with MFA](#)
- [Permissions Boundaries](#)
- [Attribute-based access control \(ABAC\)](#)

Hands-on

- Lab: [IAM Permission Boundaries Delegating Role Creation](#)
- Lab: [IAM Tag Based Access Control for EC2](#)
- Lab: [Lambda Cross Account IAM Role Assumption](#)

Detection

Detection enables you to identify a potential security misconfiguration, threat, or unexpected behavior. It's an essential part of the security lifecycle and can be used to support a quality process, a legal or compliance obligation, and for threat identification and response efforts. There are different types of detection mechanisms. For example, logs from your workload can be analyzed for exploits that are being used. You should regularly review the detection mechanisms related to your workload to ensure that you are meeting internal and external policies and requirements. Automated alerting and notifications should be based on defined conditions to enable your teams or tools to investigate. These mechanisms are important reactive factors that can help your organization identify and understand the scope of anomalous activity.

In AWS, there are a number of approaches you can use when addressing detective mechanisms. The following sections describe how to use these approaches:

- Configure
- Investigate

Configure

Configure service and application logging: A foundational practice is to establish a set of detection mechanisms at the account level. This base set of mechanisms is aimed at recording and detecting a wide range of actions on all resources in your account. They allow you to build out a comprehensive detective capability with options that include automated remediation, and partner integrations to add functionality.

In AWS, services in this base set include:

- [AWS CloudTrail](#) provides event history of your AWS account activity, including actions taken through the AWS Management Console, AWS SDKs, command line tools, and other AWS services.
- [AWS Config](#) monitors and records your AWS resource configurations and allows you to automate the evaluation and remediation against desired configurations.
- [Amazon GuardDuty](#) is a threat detection service that continuously monitors for malicious activity and unauthorized behavior to protect your AWS accounts and workloads.

- [AWS Security Hub](#) provides a single place that aggregates, organizes, and prioritizes your security alerts, or findings, from multiple AWS services and optional third-party products to give you a comprehensive view of security alerts and compliance status.

Building on the foundation at the account level, many core AWS services, for example Amazon [Virtual Private Cloud \(VPC\)](#), provide service-level logging features. [VPC Flow Logs](#) enable you to capture information about the IP traffic going to and from network interfaces that can provide valuable insight into connectivity history, and trigger automated actions based on anomalous behavior.

For EC2 instances and application-based logging that doesn't originate from AWS services, logs can be stored and analyzed using [Amazon CloudWatch Logs](#). An [agent](#) collects the logs from the operating system and the applications that are running and automatically stores them. Once the logs are available in CloudWatch Logs, you can [process them in real-time](#), or dive into analysis using [Insights](#).

Equally important to collecting and aggregating logs is the ability to extract meaningful insight from the great volumes of log and event data generated by complex architectures. See the [Monitoring](#) section of [The Reliability Pillar](#) whitepaper for more detail. Logs can themselves contain data that is considered sensitive—either when application data has erroneously found its way into log files that the CloudWatch Logs agent is capturing, or when cross-region logging is configured for log aggregation and there are legislative considerations about shipping certain kinds of information across borders.

One approach is to use Lambda functions, triggered on events when logs are delivered, to filter and redact log data before forwarding into a central logging location, such as an S3 bucket. The unredacted logs can be retained in a local bucket until a “reasonable time” has passed (as determined by legislation and your legal team), at which point an S3 lifecycle rule can automatically delete them. Logs can further be protected in Amazon S3 by using [S3 Object Lock](#), where you can store objects using a write-once-read-many (WORM) model.

Analyze logs, findings, and metrics centrally: Security operations teams rely on the collection of logs and the use of search tools to discover potential events of interest, which might indicate unauthorized activity or unintentional change. However, simply analyzing collected data and manually processing information is insufficient to keep up with the volume of information flowing from complex architectures. Analysis and reporting alone don't facilitate the assignment of the right resources to work an event in a timely fashion.

A best practice for building a mature security operations team is to deeply integrate the flow of security events and findings into a notification and workflow system such as a ticketing system, a bug/issue system, or other security information and event management (SIEM) system. This takes the workflow out of email and static reports, and allows you to route, escalate, and manage events or findings. Many organizations are also integrating security alerts into their chat/collaboration and developer productivity platforms. For organizations embarking on automation, an API-driven, low-latency ticketing system offers considerable flexibility when planning “what to automate first”.

This best practice applies not only to security events generated from log messages depicting user activity or network events, but also from changes detected in the infrastructure itself. The ability to detect change, determine whether a change was appropriate, and then route that information to the correct remediation workflow is essential in maintaining and validating a secure architecture, in the context of changes where the nature of their undesirability is sufficiently subtle that their execution cannot currently be prevented with a combination of IAM and Organizations configuration.

GuardDuty and Security Hub provide aggregation, deduplication, and analysis mechanisms for log records that are also made available to you via other AWS services. Specifically, GuardDuty ingests, aggregates, and analyses information from the VPC DNS service and information which you can otherwise see via CloudTrail and VPC Flow Logs. Security Hub can ingest, aggregate, and analyze output from GuardDuty, AWS Config, Amazon Inspector, Macie, AWS Firewall Manager, and a significant number of third-party security products available in the AWS Marketplace, and if built accordingly, your own code. Both GuardDuty and Security Hub have a Master-Member model that can aggregate findings and insights across multiple accounts, and Security Hub is often used by customers who have an on-premises SIEM as an AWS-side log and alert preprocessor and aggregator from which they can then ingest Amazon EventBridge via a Lambda-based processor and forwarder.

Resources

Refer to the following resources to learn more about current AWS recommendations for capturing and analyzing logs.

Videos

- [Threat management in the cloud: Amazon GuardDuty & AWS Security Hub](#)
- [Centrally Monitoring Resource Configuration & Compliance](#)

Documentation

- [Setting up Amazon GuardDuty](#)
- [AWS Security Hub](#)
- [Getting started: Amazon CloudWatch Logs](#)
- [Amazon EventBridge](#)
- [Configuring Athena to analyze CloudTrail logs](#)
- [Amazon CloudWatch](#)
- [AWS Config](#)
- [Creating a trail in CloudTrail](#)
- [Centralize logging solution](#)

Hands-on

- Lab: [Enable Security Hub](#)
- Lab: [Automated Deployment of Detective Controls](#)
- Lab: [Amazon GuardDuty hands on](#)

Investigate

Implement actionable security events: For each detective mechanism you have, you should also have a process, in the form of a [runbook](#) or [playbook](#), to investigate. For example, when you enable [Amazon GuardDuty](#), it generates different [findings](#). You should have a runbook entry for each finding type, for example, if a [trojan](#) is discovered, your runbook has simple instructions that instruct someone to investigate and remediate.

Automate response to events: In AWS, investigating events of interest and information on potentially unexpected changes into an automated workflow can be achieved using [Amazon EventBridge](#). This service provides a scalable rules engine designed to broker both native AWS event formats (such as CloudTrail events), as well as custom events you can generate from your application. Amazon EventBridge also allows you to route events to a workflow system for those building incident response systems (Step Functions), or to a central Security Account, or to a bucket for further analysis.

Detecting change and routing this information to the correct workflow can also be accomplished using AWS Config rules. AWS Config detects changes to in-scope services (though with higher latency than Amazon EventBridge) and generates events that can be parsed using AWS Config rules for rollback, enforcement of compliance policy, and forwarding of information to systems, such as change management platforms and operational ticketing systems. As well as writing your own Lambda functions to respond to AWS Config events, you can also take advantage of the [AWS Config Rules Development Kit](#), and a [library of open source](#) AWS Config Rules.

Resources

Refer to the following resources to learn more about current AWS best practices for integrating auditing controls with notification and workflow.

Videos

- [Amazon Detective](#)
- [Remediating Amazon GuardDuty and AWS Security Hub Findings](#)
- [Best Practices for Managing Security Operations on AWS](#)
- [Achieving Continuous Compliance using AWS Config](#)

Documentation

- [Amazon Detective](#)
- [Amazon EventBridge](#)
- [AWS Config Rules](#)
- [AWS Config Rules Repository \(open source\)](#)
- [AWS Config Rules Development Kit](#)

Hands-on

- Solution: [Real-Time Insights on AWS Account Activity](#)
- Solution: [Centralized Logging](#)

Infrastructure Protection

Infrastructure protection encompasses control methodologies, such as defense in depth, that are necessary to meet best practices and organizational or regulatory

obligations. Use of these methodologies is critical for successful, ongoing operations in the cloud.

Infrastructure protection is a key part of an information security program. It ensures that systems and services within your workload are protected against unintended and unauthorized access, and potential vulnerabilities. For example, you'll define trust boundaries (for example, network and account boundaries), system security configuration and maintenance (for example, hardening, minimization and patching), operating system authentication and authorizations (for example, users, keys, and access levels), and other appropriate policy-enforcement points (for example, web application firewalls and/or API gateways).

In AWS, there are a number of approaches to infrastructure protection. The following sections describe how to use these approaches:

- Protecting networks
- Protecting compute

Protecting Networks

The careful planning and management of your network design forms the foundation of how you provide isolation and boundaries for resources within your workload. Because many resources in your workload operate in a VPC and inherit the security properties, it's critical that the design is supported with inspection and protection mechanisms backed by automation. Likewise, for workloads that operate outside a VPC, using purely edge services and/or serverless, the best practices apply in a more simplified approach. Refer to the [AWS Well-Architected Serverless Applications Lens](#) for specific guidance on serverless security.

Create network layers: Components such as EC2 instances, RDS database clusters, and Lambda functions that share reachability requirements can be segmented into layers formed by subnets. For example, an RDS database cluster in a VPC with no need for internet access should be placed in subnets with no route to or from the internet. This layered approach for the controls mitigates the impact of a single layer misconfiguration, which could allow unintended access. For AWS Lambda, you can run your functions in your VPC to take advantage of VPC-based controls.

For network connectivity that can include thousands of VPCs, AWS accounts, and on-premises networks, you should use [AWS Transit Gateway](#). It acts as a hub that controls how traffic is routed among all the connected networks, which act like spokes. Traffic

between an Amazon VPC and AWS Transit Gateway remains on the AWS private network, which reduces external threat vectors such as distributed denial of service (DDoS) attacks and common exploits, such as SQL injection, cross-site scripting, cross-site request forgery, or abuse of broken authentication code. AWS Transit Gateway inter-region peering also encrypts inter-region traffic with no single point of failure or bandwidth bottleneck.

Control traffic at all layers: When architecting your network topology, you should examine the connectivity requirements of each component. For example, if a component requires internet accessibility (inbound and outbound), connectivity to VPCs, edge services, and external data centers.

A VPC allows you to define your network topology that spans an AWS Region with a private IPv4 address range that you set, or an IPv6 address range AWS selects. You should apply multiple controls with a defense in depth approach for both inbound and outbound traffic, including the use of security groups (stateful inspection firewall), Network ACLs, subnets, and route tables. Within a VPC, you can create subnets in an Availability Zone. Each subnet can have an associated route table that defines routing rules for managing the paths that traffic takes within the subnet. You can define an internet routable subnet by having a route that goes to an internet or NAT gateway attached to the VPC, or through another VPC.

When an instance, RDS database, or other service is launched within a VPC, it has its own security group per network interface. This firewall is outside the operating system layer and can be used to define rules for allowed inbound and outbound traffic. You can also define relationships between security groups. For example, instances within a database tier security group only accept traffic from instances within the application tier, by reference to the security groups applied to the instances involved. Unless you are using non-TCP protocols, it shouldn't be necessary to have an EC2 instance directly accessible by the internet (even with ports restricted by security groups) without a load balancer, or [CloudFront](#). This helps protect it from unintended access through an operating system or application issue. A subnet can also have a network ACL attached to it, which acts as a stateless firewall. You should configure the network ACL to narrow the scope of traffic allowed between layers, note that you need to define both inbound and outbound rules.

While some AWS services require components to access the internet to make API calls (this being where AWS API [endpoints are located](#)), others use [endpoints](#) within your VPCs. Many AWS services including Amazon S3 and DynamoDB support VPC endpoints, and this technology has been generalized in AWS PrivateLink. For VPC

assets that need to make outbound connections to the internet, these can be made outbound only (one-way) through an AWS managed NAT gateway, outbound only internet gateway, or web proxies that you create and manage.

Implement inspection and protection: Inspect and filter your traffic at each layer. For components transacting over HTTP-based protocols, a web application firewall can help protect from common attacks. [AWS WAF](#) is a web application firewall that lets you monitor and block HTTP(s) requests that match your configurable rules that are forwarded to an Amazon API Gateway API, Amazon CloudFront, or an Application Load Balancer. To get started with AWS WAF, you can use [AWS Managed Rules](#) in combination with your own, or use existing [partner integrations](#).

For managing AWS WAF, AWS Shield Advanced protections, and Amazon VPC security groups across AWS Organizations, you can use AWS Firewall Manager. It allows you to centrally configure and manage firewall rules across your accounts and applications, making it easier to scale enforcement of common rules. It also enables you to rapidly respond to attacks, using [AWS Shield Advanced](#), or [solutions](#) that can automatically block unwanted requests to your web applications.

Automate network protection: Automate protection mechanisms to provide a self-defending network based on threat intelligence and anomaly detection. For example, intrusion detection and prevention tools that can adapt to current threats and reduce their impact. A web application firewall is an example of where you can automate network protection, for example, by using the [AWS WAF Security Automations solution \(https://github.com/awslabs/aws-waf-security-automations\)](#) to automatically block requests originating from IP addresses associated with known threat actors.

Resources

Refer to the following resources to learn more about AWS best practices for protecting networks.

Video

- [AWS Transit Gateway reference architectures for many VPCs](#)
- [Application Acceleration and Protection with Amazon CloudFront, AWS WAF, and AWS Shield](#)
- [DDoS Attack Detection at Scale](#)

Documentation

- [Amazon VPC Documentation](#)
- [Getting started with AWS WAF](#)
- [Network Access Control Lists](#)
- [Security Groups for Your VPC](#)
- [Recommended Network ACL Rules for Your VPC](#)
- [AWS Firewall Manager](#)
- [AWS PrivateLink](#)
- [VPC Endpoints](#)
- [Amazon Inspector](#)

Hands-on

- Lab: [Automated Deployment of VPC](#)
- Lab: [Automated Deployment of Web Application Firewall](#)

Protecting Compute

Perform vulnerability management: Frequently scan and patch for vulnerabilities in your code, dependencies, and in your infrastructure to help protect against new threats.

Using a build and deployment pipeline, you can automate many parts of vulnerability management:

- Using third-party static code analysis tools to identify common security issues such as unchecked function input bounds, as well as more recent CVEs. You can use [Amazon CodeGuru](#) for languages supported.
- Using third-party dependency checking tools to determine whether libraries your code links against are the latest versions, are themselves free of CVEs, and have licensing conditions that meet your software policy requirements.

- Using Amazon Inspector, you can perform configuration assessments against your instances for known common vulnerabilities and exposures (CVEs), assess against security benchmarks, and fully automate the notification of defects. Amazon Inspector runs on production instances or in a build pipeline, and it notifies developers and engineers when findings are present. You can access findings programmatically and direct your team to backlogs and bug-tracking systems. [EC2 Image Builder](#) can be used to maintain server images (AMIs) with automated patching, AWS-provided security policy enforcement, and other customizations.
- When using containers implement [ECR Image Scanning](#) in your build pipeline and on a regular basis against your image repository to look for CVEs in your containers.
- While Amazon Inspector and other tools are effective at identifying configurations and any CVEs that are present, other methods are required to test your workload at the application level. [Fuzzing](#) is a well-known method of finding bugs using automation to inject malformed data into input fields and other areas of your application.

A number of these functions can be performed using AWS services, products in the AWS Marketplace, or open-source tooling.

Reduce attack surface: Reduce your attack surface by hardening operating systems, minimizing components, libraries, and externally consumable services in use. To reduce your attack surface, you need a threat model to identify the entry points and potential threats that could be encountered. A common practice in reducing attack surface is to start at reducing unused components, whether they are operating system packages, applications, etc (for EC2 based workloads) or external software modules in your code (for all workloads). Many hardening and security configuration guides exist for common operating systems and server software, for example from the [Center for Internet Security](#) that you can use as a starting point and iterate.

Enable people to perform actions at a distance: Removing the ability for interactive access reduces the risk of human error, and the potential for manual configuration or management. For example, use a change management workflow to manage EC2 instances using tools such as AWS Systems Manager instead of allowing direct access, or via a bastion host. AWS Systems Manager can automate a variety of maintenance and deployment tasks, using features including [automation workflows](#), [documents](#) (playbooks), and the [run command](#). AWS CloudFormation stacks build from pipelines

and can automate your infrastructure deployment and management tasks without using the AWS Management Console or APIs directly.

Implement managed services: Implement services that manage resources, such as Amazon RDS, AWS Lambda, and Amazon ECS, to reduce your security maintenance tasks as part of the shared responsibility model. For example, Amazon RDS helps you set up, operate, and scale a relational database, automates administration tasks such as hardware provisioning, database setup, patching, and backups. This means you have more free time to focus on securing your application in other ways described in the AWS Well-Architected Framework. AWS Lambda lets you run code without provisioning or managing servers, so you only need to focus on the connectivity, invocation, and security at the code level—not the infrastructure or operating system.

Validate software integrity: Implement mechanisms (e.g. code signing) to validate that the software, code and libraries used in the workload are from trusted sources and have not been tampered with. For example, you should verify the code signing certificate of binaries and scripts to confirm the author, and ensure it has not been tampered with since created by the author. Additionally, a checksum of software that you download, compared to that of the checksum from the provider, can help ensure it has not been tampered with.

Automate compute protection: Automate your protective compute mechanisms including vulnerability management, reduction in attack surface, and management of resources. The automation will help you invest time in securing other aspects of your workload, and reduce the risk of human error.

Resources

Refer to the following resources to learn more about AWS best practices for protecting compute.

Video

- [Security best practices for the Amazon EC2 instance metadata service](#)
- [Securing Your Block Storage on AWS](#)
- [Securing Serverless and Container Services](#)
- [Running high-security workloads on Amazon EKS](#)
- [Architecting Security through Policy Guardrails in Amazon EKS](#)

Documentation

- [Security Overview of AWS Lambda](#)
- [Security in Amazon EC2](#)
- [AWS Systems Manager](#)
- [Amazon Inspector](#)
- [Writing your own AWS Systems Manager documents](#)
- [Replacing a Bastion Host with Amazon EC2 Systems Manager](#)

Hands-on

- Lab: [Automated Deployment of EC2 Web Application](#)

Data Protection

Before architecting any workload, foundational practices that influence security should be in place. For example, data classification provides a way to categorize data based on levels of sensitivity, and encryption protects data by way of rendering it unintelligible to unauthorized access. These methods are important because they support objectives such as preventing mishandling or complying with regulatory obligations.

In AWS, there are a number of different approaches you can use when addressing data protection. The following section describes how to use these approaches:

- Data classification
- Protecting data at rest
- Protecting data in transit

Data Classification

Data classification provides a way to categorize organizational data based on criticality and sensitivity in order to help you determine appropriate protection and retention controls.

Identify the data within your workload: You need to understand the type and classification of data your workload is processing, the associated business processes, data owner, applicable legal and compliance requirements, where it's stored, and the resulting controls that are needed to be enforced. This may include classifications to indicate if the data is intended to be publicly available, if the data is internal use only such as customer personally identifiable information (PII), or if the data is for more restricted access such as intellectual property, legally privileged or marked sensitive, and more. By carefully managing an appropriate data classification system, along with each workload's level of protection requirements, you can map the controls and level of access/protection appropriate for the data. For example, public content is available for anyone to access, but important content is encrypted and stored in a protected manner that requires authorized access to a key for decrypting the content.

Define data protection controls: By using resource tags, separate AWS accounts per sensitivity (and potentially also per caveat / enclave / community of interest), IAM policies, Organizations SCPs, AWS KMS, and AWS CloudHSM, you can define and implement your policies for data classification and protection with encryption. For example, if you have a project with S3 buckets that contain highly critical data or EC2

instances that process confidential data, they can be tagged with a “Project=ABC” tag. Only your immediate team knows what the project code means, and it provides a way to use attribute-based access control. You can define levels of access to the AWS KMS encryption keys through key policies and grants to ensure that only appropriate services have access to the sensitive content through a secure mechanism. If you are making authorization decisions based on tags you should make sure that the permissions on the tags are defined appropriately using tag policies in AWS Organizations.

Define data lifecycle management: Your defined lifecycle strategy should be based on sensitivity level as well as legal and organization requirements. Aspects including the duration for which you retain data, data destruction processes, data access management, data transformation, and data sharing should be considered. When choosing a data classification methodology, balance usability versus access. You should also accommodate the multiple levels of access and nuances for implementing a secure, but still usable, approach for each level. Always use a defense in depth approach and reduce human access to data and mechanisms for transforming, deleting, or copying data. For example, require users to strongly authenticate to an application, and give the application, rather than the users, the requisite access permission to perform “action at a distance.” In addition, ensure that users come from a trusted network path and require access to the decryption keys. Use tools, such as dashboards and automated reporting, to give users information from the data rather than giving them direct access to the data.

Automate identification and classification: Automating the identification and classification of data can help you implement the correct controls. Using automation for this instead of direct access from a person reduces the risk of human error and exposure. You should evaluate using a tool, such as [Amazon Macie](#), that uses machine learning to automatically discover, classify, and protect sensitive data in AWS. Amazon Macie recognizes sensitive data, such as personally identifiable information (PII) or intellectual property, and provides you with dashboards and alerts that give visibility into how this data is being accessed or moved.

Resources

Refer to the following resources to learn more about data classification.

Documentation

- [Data Classification Whitepaper](#)
- [Tagging Your Amazon EC2 Resources](#)

- [Amazon S3 Object Tagging](#)

Protecting Data at Rest

Data at rest represents any data that you persist in non-volatile storage for any duration in your workload. This includes block storage, object storage, databases, archives, IoT devices, and any other storage medium on which data is persisted. Protecting your data at rest reduces the risk of unauthorized access, when encryption and appropriate access controls are implemented.

Encryption and tokenization are two important but distinct data protection schemes.

Tokenization is a process that allows you to define a token to represent an otherwise sensitive piece of information (for example, a token to represent a customer's credit card number). A token must be meaningless on its own, and must not be derived from the data it is tokenizing—therefore, a cryptographic digest is not usable as a token. By carefully planning your tokenization approach, you can provide additional protection for your content, and you can ensure that you meet your compliance requirements. For example, you can reduce the compliance scope of a credit card processing system if you leverage a token instead of a credit card number.

Encryption is a way of transforming content in a manner that makes it unreadable without a secret key necessary to decrypt the content back into plaintext. Both tokenization and encryption can be used to secure and protect information as appropriate. Further, masking is a technique that allows part of a piece of data to be redacted to a point where the remaining data is not considered sensitive. For example, PCI-DSS allows the last four digits of a card number to be retained outside the compliance scope boundary for indexing.

Implement secure key management: By defining an encryption approach that includes the storage, rotation, and access control of keys, you can help provide protection for your content against unauthorized users and against unnecessary exposure to authorized users. AWS KMS helps you manage encryption keys and [integrates with many AWS services](#). This service provides durable, secure, and redundant storage for your master keys. You can define your key aliases as well as key-level policies. The policies help you define key administrators as well as key users. Additionally, AWS CloudHSM is a cloud-based hardware security module (HSM) that enables you to easily generate and use your own encryption keys in the AWS Cloud. It helps you meet corporate, contractual, and regulatory compliance requirements for data security by using FIPS 140-2 Level 3 validated HSMs.

Enforce encryption at rest: You should ensure that the only way to store data is by using encryption. AWS KMS integrates seamlessly with many AWS services to make it easier for you to encrypt all your data at rest. For example, in Amazon S3 you can set [default encryption](#) on a bucket so that all new objects are automatically encrypted. Additionally, Amazon EC2 supports the enforcement of encryption by [setting a default encryption option](#) for an entire Region.

Enforce access control: Different controls including access (using least privilege), backups (see Reliability whitepaper), isolation, and versioning can all help protect your data at rest. Access to your data should be audited using detective mechanisms covered earlier in this paper including CloudTrail, and service level log, such as S3 access logs. You should inventory what data is publicly accessible, and plan for how you can reduce the amount of data available over time. Amazon S3 Glacier Vault Lock and S3 Object Lock are capabilities providing mandatory access control—once a vault policy is locked with the compliance option, not even the root user can change it until the lock expires. The mechanism meets the Books and Records Management requirements of the SEC, CFTC, and FINRA. For more details, see [this whitepaper](#).

Audit the use of encryption keys: Ensure that you understand and audit the use of encryption keys to validate that the access control mechanisms on the keys are appropriately implemented. For example, any AWS service using an AWS KMS key logs each use in AWS CloudTrail. You can then query AWS CloudTrail, by using a tool such as Amazon CloudWatch Insights, to ensure that all uses of your keys are valid.

Use mechanisms to keep people away from data: Keep all users away from directly accessing sensitive data and systems under normal operational circumstances. For example, use a change management workflow to manage EC2 instances using tools instead of allowing direct access or a bastion host. This can be achieved using [AWS Systems Manager Automation](#), which uses [automation documents](#) that contain steps you use to perform tasks. These documents can be stored in source control, be peer reviewed before running, and tested thoroughly to minimize risk compared to shell access. Business users could have a dashboard instead of direct access to a data store to run queries. Where CI/CD pipelines are not used, determine which controls and processes are required to adequately provide a normally disabled break-glass access mechanism.

Automate data at rest protection: Use automated tools to validate and enforce data at rest controls continuously, for example, verify that there are only encrypted storage resources. You can [automate validation that all EBS volumes are encrypted](#) using [AWS Config Rules](#). [AWS Security Hub](#) can also verify a number of different controls through

automated checks against security standards. Additionally, your AWS Config Rules can automatically [remediate noncompliant resources](#).

Resources

Refer to the following resources to learn more about AWS best practices for protecting data at rest.

Video

- [How Encryption Works in AWS](#)
- [Securing Your Block Storage on AWS](#)
- [Achieving security goals with AWS CloudHSM](#)
- [Best Practices for Implementing AWS Key Management Service](#)
- [A Deep Dive into AWS Encryption Services](#)

Documentation

- [Protecting Amazon S3 Data Using Encryption](#)
- [Amazon EBS Encryption](#)
- [Encrypting Amazon RDS Resources](#)
- [Protecting Data Using Encryption](#)
- [How AWS services use AWS KMS](#)
- [Amazon EBS Encryption](#)
- [AWS Key Management Service](#)
- [AWS CloudHSM](#)
- [AWS KMS Cryptographic Details Whitepaper](#)
- [Using Key Policies in AWS KMS](#)
- [Using Bucket Policies and User Policies](#)
- [AWS Crypto Tools](#)

Protecting Data in Transit

Data in transit is any data that is sent from one system to another. This includes communication between resources within your workload as well as communication between other services and your end users. By providing the appropriate level of protection for your data in transit, you protect the confidentiality and integrity of your workload's data.

Implement secure key and certificate management: Store encryption keys and certificates securely and rotate them at appropriate time intervals with strict access control. The best way to accomplish this is to use a managed service, such as [AWS Certificate Manager \(ACM\)](#). It lets you easily provision, manage, and deploy public and private Transport Layer Security (TLS) certificates for use with AWS services and your internal connected resources. TLS certificates are used to secure network communications and establish the identity of websites over the internet as well as resources on private networks. ACM integrates with AWS resources, such as Elastic Load Balancers, Amazon CloudFront distributions, and APIs on API Gateway, also handling automatic certificate renewals. If you use ACM to deploy a private root CA, both certificates and private keys can be provided by it for use in EC2 instances, containers, etc.

Enforce encryption in transit: Enforce your defined encryption requirements based on appropriate standards and recommendations to help you meet your organizational, legal, and compliance requirements. AWS services provide HTTPS endpoints using TLS for communication, thus providing encryption in transit when communicating with the AWS APIs. Insecure protocols, such as HTTP, can be audited and blocked in a VPC through the use of security groups. HTTP requests can also be [automatically redirected to HTTPS](#) in Amazon CloudFront or on an [Application Load Balancer](#). You have full control over your computing resources to implement encryption in transit across your services. Additionally, you can use VPN connectivity into your VPC from an external network to facilitate encryption of traffic. Third-party solutions are available in the AWS Marketplace, if you have special requirements.

Authenticate network communications: Using network protocols that support authentication allows for trust to be established between the parties. This adds to the encryption used in the protocol to reduce the risk of communications being altered or intercepted. Common protocols that implement authentication include Transport Layer Security (TLS), which is used in many AWS services, and IPsec, which is used in [AWS Virtual Private Network \(AWS VPN\)](#).

Automate detection of unintended data access: Use tools such as Amazon GuardDuty to automatically detect attempts to move data outside of defined boundaries based on data classification level, for example, to detect a trojan that is copying data to an unknown or untrusted network using the DNS protocol. In addition to Amazon GuardDuty, [Amazon VPC Flow Logs](#), which capture network traffic information, can be used with Amazon EventBridge to trigger detection of abnormal connections—both successful and denied. [S3 Access Analyzer](#) can help assess what data is accessible to who in your S3 buckets.

Resources

Refer to the following resources to learn more about AWS best practices for protecting data in transit.

Video

- [How can I add certificates for websites to the ELB using AWS Certificate Manager](#)
- [Deep Dive on AWS Certificate Manager Private CA](#)

Documentation

- [AWS Certificate Manager](#)
- [HTTPS Listeners for Your Application Load Balancer](#)
- [AWS VPN](#)
- [API Gateway Edge-Optimized](#)

Incident Response

Even with extremely mature preventive and detective controls, your organization should still implement mechanisms to respond to and mitigate the potential impact of security incidents. Your preparation strongly affects the ability of your teams to operate effectively during an incident, to isolate and contain issues, and to restore operations to a known good state. Putting in place the tools and access ahead of a security incident, then routinely practicing incident response through game days, helps ensure that you can recover while minimizing business disruption.

Design Goals of Cloud Response

Although the general processes and mechanisms of incident response, such as those defined in the [NIST SP 800-61 Computer Security Incident Handling Guide](#), remain true, we encourage you to evaluate these specific design goals that are relevant to responding to security incidents in a cloud environment:

- **Establish response objectives:** Work with your stakeholders, legal counsel, and organizational leadership to determine the goal of responding to an incident. Some common goals include containing and mitigating the issue, recovering the affected resources, preserving data for forensics, and attribution.
- **Document plans:** Create plans to help you respond to, communicate during, and recover from an incident.
- **Respond using the cloud:** Implement your response patterns where the event and data occurs.
- **Know what you have and what you need:** Preserve logs, snapshots, and other evidence by copying them to a centralized security cloud account. Use tags, metadata, and mechanisms that enforce retention policies. For example, you might choose to use the Linux `dd` command or a Windows equivalent to make a complete copy of the data for investigative purposes.
- **Use redeployment mechanisms:** If a security anomaly can be attributed to a misconfiguration, the remediation might be as simple as removing the variance by redeploying the resources with the proper configuration. When possible, make your response mechanisms safe to execute more than once and in environments in an unknown state.

- **Automate where possible:** As you see issues or incidents repeat, build mechanisms that programmatically triage and respond to common situations. Use human responses for unique, new, and sensitive incidents.
- **Choose scalable solutions:** Strive to match the scalability of your organization's approach to cloud computing, and reduce the time between detection and response.
- **Learn and improve your process:** When you identify gaps in your process, tools, or people, implement plans to fix them. Simulations are safe methods to find gaps and improve processes.

In AWS, there are a number of different approaches you can use when addressing incident response. The following section describes how to use these approaches:

- **Educate** your security operations and incident response staff about cloud technologies and how your organization intends to use them.
- **Prepare** your incident response team to detect and respond to incidents in the cloud, enable detective capabilities, and ensure appropriate access to the necessary tools and cloud services. Additionally, prepare the necessary runbooks, both manual and automated, to ensure reliable and consistent responses. Work with other teams to establish expected baseline operations, and use that knowledge to identify deviations from those normal operations.
- **Simulate** both expected and unexpected security events within your cloud environment to understand the effectiveness of your preparation.
- **Iterate** on the outcome of your simulation to improve the scale of your response posture, reduce time to value, and further reduce risk.

Educate

Automated processes enable organizations to spend more time focusing on measures to increase the security of their workloads. Automated incident response also makes humans available to correlate events, practice simulations, devise new response procedures, perform research, develop new skills, and test or build new tools. Despite increased automation, your team, specialists, and responders within a security organization still require continuous education.

Beyond general cloud experience, you need to significantly invest in your people to be successful. Your organization can benefit by providing additional training to your staff to learn programming skills, development processes (including version control systems

and deployment practices), and infrastructure automation. The best way to learn is hands-on, through running incident response game days. This allows for experts in your team to hone the tools and techniques while teaching others.

Prepare

During an incident, your incident response teams must have access to various tools and the workload resources involved in the incident. Make sure that your teams have appropriate pre-provisioned access to perform their duties before an event occurs. All tools, access, and plans should be documented and tested before an event occurs to make sure that they can provide a timely response.

Identify key personnel and external resources: When you define your approach to incident response in the cloud, in unison with other teams (such as your legal counsel, leadership, business stakeholders, AWS Support Services, and others), you must identify key personnel, stakeholders, and relevant contacts. To reduce dependency and decrease response time, make sure that your team, specialist security teams, and responders are educated about the services that you use and have opportunities to practice hands-on.

We encourage you to identify external AWS security partners that can provide you with outside expertise and a different perspective to augment your response capabilities. Your trusted security partners can help you identify potential risks or threats that you might not be familiar with.

Develop incident management plans: Create plans to help you respond to, communicate during, and recover from an incident. For example, you can start at incident response plan with the most likely scenarios for your workload and organization. Include how you would communicate and escalate both internally and externally. Create incident response plans in the form of [playbooks](#), starting with the most likely scenarios for your workload and organization. These might be events that are currently generated. If you need a starting place, you should look at [AWS Trusted Advisor](#) and [Amazon GuardDuty findings](#). Use a simple format such as markdown so it's easily maintained but ensure that important commands or code snippets are included so they can be executed without having to lookup other documentation.

Start simple and iterate. Work closely with your security experts and partners to identify the tasks required to ensure that the processes are possible. Define the manual descriptions of the processes you perform. After this, test the processes and iterate on the runbook pattern to improve the core logic of your response. Determine what the exceptions are, and what the alternative resolutions are for those scenarios. For

example, in a development environment, you might want to terminate a misconfigured Amazon EC2 instance. But, if the same event occurred in a production environment, instead of terminating the instance, you might stop the instance and verify with stakeholders that critical data will not be lost and that termination is acceptable. Include how you would communicate and escalate both internally and externally. When you are comfortable with the manual response to the process, automate it to reduce the time to resolution.

Pre-provision access: Ensure that incident responders have the correct access pre-provisioned into AWS and other relevant systems to reduce the time for investigation through to recovery. Determining how to get access for the right people during an incident delays the time it takes to respond, and can introduce other security weaknesses if access is shared or not properly provisioned while under pressure. You must know what level of access your team members require (for example, what kinds of actions they are likely to take) and you must provision access in advance. Access in the form of roles or users created specifically to respond to a security incident are often privileged in order to provide sufficient access. Therefore, use of these user accounts should be restricted, they should not be used for daily activities, and usage alerted on.

Pre-deploy tools: Ensure that security personnel have the right tools pre-deployed into AWS to reduce the time for investigation through to recovery.

To automate security engineering and operations functions, you can use a comprehensive set of APIs and tools from AWS. You can fully automate identity management, network security, data protection, and monitoring capabilities and deliver them using popular software development methods that you already have in place. When you build security automation, your system can monitor, review, and initiate a response, rather than having people monitor your security position and manually react to events.

If your incident response teams continue to respond to alerts in the same way, they risk alert fatigue. Over time, the team can become desensitized to alerts and can either make mistakes handling ordinary situations or miss unusual alerts. Automation helps avoid alert fatigue by using functions that process the repetitive and ordinary alerts, leaving humans to handle the sensitive and unique incidents.

You can improve manual processes by programmatically automating steps in the process. After you define the remediation pattern to an event, you can decompose that pattern into actionable logic, and write the code to perform that logic. Responders can then execute that code to remediate the issue. Over time, you can automate more and more steps, and ultimately automatically handle whole classes of common incidents.

For tools that execute within the operating system of your EC2 instance, you should evaluate using the AWS Systems Manager Run Command, which enables you to remotely and securely administrate instances using an agent that you install on your Amazon EC2 instance operating system. It requires the AWS Systems Manager Agent (SSM Agent), which is installed by default on many Amazon Machine Images (AMIs). Be aware, though, that once an instance has been compromised, no responses from tools or agents running on it should be considered trustworthy.

Prepare forensic capabilities: Identify and prepare forensic investigation capabilities that are suitable, including external specialists, tools, and automation. Some of your incident response activities might include analyzing disk images, file systems, RAM dumps, or other artifacts that are involved in an incident. Build a customized forensic workstation that they can use to mount copies of any affected data volumes. As forensic investigation techniques require specialist training, you might need to engage external specialists.

Simulate

Run game days: Game days, also known as simulations or exercises, are internal events that provide a structured opportunity to practice your incident management plans and procedures during a realistic scenario. Game days are fundamentally about being prepared and iteratively improving your response capabilities. Some of the reasons you might find value in performing game day activities include:

- Validating readiness
- Developing confidence – learning from simulations and training staff
- Following compliance or contractual obligations
- Generating artifacts for accreditation
- Being agile – incremental improvement
- Becoming faster and improving tools
- Refining communication and escalation
- Developing comfort with the rare and the unexpected

For these reasons, the value derived from participating in a SIRS activity increases an organization's effectiveness during stressful events. Developing a SIRS activity that is both realistic and beneficial can be a difficult exercise. Although testing your procedures or automation that handles well-understood events has certain advantages, it is just as

valuable to participate in creative SIRS activities to test yourself against the unexpected and continuously improve.

Iterate

Automate containment and recovery capability: Automate containment and recovery of an incident to reduce response times and organizational impact.

Once you create and practice the processes and tools from your playbooks, you can deconstruct the logic into a code-based solution, which can be used as a tool by many responders to automate the response and remove variance or guess-work by your responders. This can speed up the lifecycle of a response. The next goal is to enable this code to be fully automated by being invoked by the alerts or events themselves, rather than by a human responder, to create an event-driven response.

With an event-driven response system, a detective mechanism triggers a responsive mechanism to automatically remediate the event. You can use event-driven response capabilities to reduce the time-to-value between detective mechanisms and responsive mechanisms. To create this event-driven architecture, you can use AWS Lambda, which is a serverless compute service that runs your code in response to events and automatically manages the underlying compute resources for you. For example, assume that you have an AWS account with the AWS CloudTrail service enabled. If AWS CloudTrail is ever disabled (through the `cloudtrail:StopLogging` API call), you can use Amazon EventBridge to monitor for the specific `cloudtrail:StopLogging` event, and invoke an AWS Lambda function to call `cloudtrail:StartLogging` to re-start logging.

Resources

Refer to the following resources to learn more about current AWS best practices for incident response.

Videos

- [Prepare for & respond to security incidents in your AWS environment](#)
- [Automating Incident Response and Forensics](#)
- [DIY guide to runbooks, incident reports, and incident response](#)

Documentation

- [AWS Incident Response Guide](#)
- [AWS Step Functions](#)
- [Amazon EventBridge](#)
- [CloudEndure Disaster Recovery](#)

Hands-on

- Lab: [Incident Response with AWS Console and CLI](#)
- Lab: [Incident Response Playbook with Jupyter - AWS IAM](#)
- Blog: [Orchestrating a security incident response with AWS Step Functions](#)

Conclusion

Security is an ongoing effort. When incidents occur, they should be treated as opportunities to improve the security of the architecture. Having strong identity controls, automating responses to security events, protecting infrastructure at multiple levels, and managing well-classified data with encryption provides defense in depth that every organization should implement. This effort is easier thanks to the programmatic functions and AWS features and services discussed in this paper.

AWS strives to help you build and operate architectures that protect information, systems, and assets while delivering business value.

Contributors

The following individuals and organizations contributed to this document:

- Ben Potter, Principal Security Lead, Well-Architected, Amazon Web Services
- Bill Shinn, Senior Principal, Office of the CISO, Amazon Web Services
- Brigid Johnson, Senior Software Development Manager, AWS Identity, Amazon Web Services
- Byron Pogson, Senior Solution Architect, Amazon Web Services
- Darran Boyd, Principal Security Solutions Architect, Financial Services, Amazon Web Services

- Dave Walker, Principal Specialist Solutions Architect, Security and Compliance, Amazon Web Services
- Paul Hawkins, Senior Security Strategist, Amazon Web Services
- Sam Elmalak, Senior Technology Leader, Amazon Web Services

Further Reading

For additional help, please consult the following sources:

- [AWS Well-Architected Framework Whitepaper](#)

Document Revisions

Date	Description
July 2020	Updated guidance on account, identity, and permissions management.
April 2020	Updated to expand advice in every area, new best practices, services and features.
July 2018	Updates to reflect new AWS services and features, and updated references.
May 2017	Updated System Security Configuration and Maintenance section to reflect new AWS services and features.
November 2016	First publication