

AWS Security Incident Response Guide

June 2019



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.

© 2019 Amazon Web Services, Inc. or its affiliates. All rights reserved.

Contents

- Introduction 1
 - Before You Begin2
 - The AWS CAF Security Perspective2
 - Foundation of Incident Response2
- Educate3
 - Shared Responsibility3
 - Incident Response in the Cloud4
 - Cloud Security Incidents6
 - Understanding Cloud Capabilities8
- Prepare – People10
 - Define Roles and Responsibilities10
 - Define Response Mechanisms11
 - Create a Receptive and Adaptive Security Culture11
 - Predicting Response12
- Prepare – Technology16
 - Prepare Access to AWS Accounts16
 - Prepare Processes19
 - Cloud Provider Support25
- Simulate27
 - Security Incident Response Simulations27
 - Simulation Steps28
 - Simulation Examples29
- Iterate30
 - Runbooks30
 - Automation31
- Incident Response Examples34

Service Domain Incidents	34
Infrastructure Domain Incidents.....	35
Conclusion	38
Contributors	38
Additional Resources	39
Media	39
Third-Party Tools.....	40
Industry References.....	40
Document Revisions.....	40
Appendix A – Examples	41
Appendix B – Cloud Capability Definitions	45
Logging and Events	45
Visibility and Alerting	47
Automation	48
Secure Storage	49
Custom	50

About This Guide

This guide presents an overview of the fundamentals of responding to security incidents within a customer's AWS Cloud environment. It focuses on an overview of cloud security and incident response concepts, and identifies cloud capabilities, services, and mechanisms that are available to customers who are responding to security issues.

Introduction

Security is the highest priority at AWS. As an AWS customer, you benefit from a data center and network architecture that is built to meet the requirements of the most security-sensitive organizations. The AWS Cloud has a *shared responsibility model*. AWS manages security *of* the cloud. You are responsible for security *in* the cloud. This means that you retain control of the security you choose to implement. You get access to hundreds of tools and features to help you to meet your security objectives. These capabilities help you establish a security baseline that meets your objectives for your applications running in the cloud.

When a deviation from your baseline does occur (such as by a misconfiguration), you might need to respond and investigate. To successfully do so, you must understand the basic concepts of security incident response within your AWS environment, as well as the issues you need to consider to prepare, educate, and train your cloud teams before security issues occur. It is important to know which controls and capabilities you can use, and can be helpful to review topical examples for resolving potential concerns, and identify remediation methods that you can use to leverage automation and improve your response speed.

Because security incident response can be a complex topic, we encourage customers to start small, develop runbooks, leverage basic capabilities, and create an initial library of incident response mechanisms to iterate from and improve upon. This initial work should include teams that are not involved with security and should include your legal department, so that you are better able to understand the impact that incident response (IR), and the choices you have made, have on your corporate goals.

This paper is intended for those in technical roles and assumes that you are familiar with the general principles of information security, have a basic understanding of incident response in your current on-premises environments, and have some familiarity with cloud services.

Before You Begin

In addition to this document, we encourage you to review the [AWS Security Best Practices](#) whitepaper and the [Security Perspective of the AWS Cloud Adoption Framework \(CAF\)](#) whitepaper. The AWS CAF provides guidance that supports coordinating between the different parts of organizations that are moving to the cloud. The CAF guidance is divided into several areas of focus that are relevant to implementing cloud-based IT systems, which we refer to as *perspectives*. The *Security Perspective* describes how to execute a security program across several workstreams, one of which focuses on incident response. This document details some of our experiences in helping customers to assess and implement successful mechanisms in that workstream.

The AWS CAF Security Perspective

The Security Perspective includes four components:

- **Directive controls** establish the governance, risk, and compliance models within which the environment operates.
- **Preventive controls** protect your workloads and mitigate threats and vulnerabilities.
- **Detective controls** provide full visibility and transparency over the operation of your deployments in AWS.
- **Responsive controls** drive remediation of potential deviations from your security baselines.

Although incident response (IR) is generally viewed under the responsive controls component, these are dependent and influenced by the other components. For example, directive and preventative security controls help establish a baseline, so you can monitor and investigate any deviations from this baseline. This approach not only eliminates noise, but it also contributes to a defensive security design.

Foundation of Incident Response

All AWS users within an organization should have a basic understanding of security incident response processes, and security staff must deeply understand how to react to security issues. Experience and education are vital to a cloud incident response program, before you handle a security event. The foundation of a successful incident response program in the cloud is to *Educate, Prepare, Simulate, and Iterate*.

To understand each of these aspects, consider the following descriptions:

- **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, enabling detective capabilities, and ensuring 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

Shared Responsibility

The responsibility for security and compliance is shared between AWS and you. This shared model relieves some of your operational burden because AWS operates, manages, and controls the components from the host operating system and virtualization layer, down to the physical security of the facilities in which the service operates.

You are responsible for managing the guest operating systems (including updates and security patches), application software, as well as configuring the AWS provided security controls, such as security groups, network access control lists, and identity and access management. You should carefully consider which services you use, because your responsibilities vary depending on the services you choose, the integration of those services in your IT environment, and applicable laws and regulations.

[Figure 1](#) shows a typical representation of the shared responsibility model as it applies to infrastructure services, such as Amazon Elastic Compute Cloud (Amazon EC2). It separates most responsibilities into two categories: security *of* the cloud (managed by AWS) and security *in* the cloud (managed by the customer). Responsibilities can change, depending on which services you use. For abstracted services, such as Amazon S3 and Amazon DynamoDB, AWS operates the infrastructure layer, the

operating system, and platforms, and customers access the endpoints to store and retrieve data. Customers are responsible for managing their data (including encryption options), classifying their assets, and using IAM tools to apply the appropriate permissions.

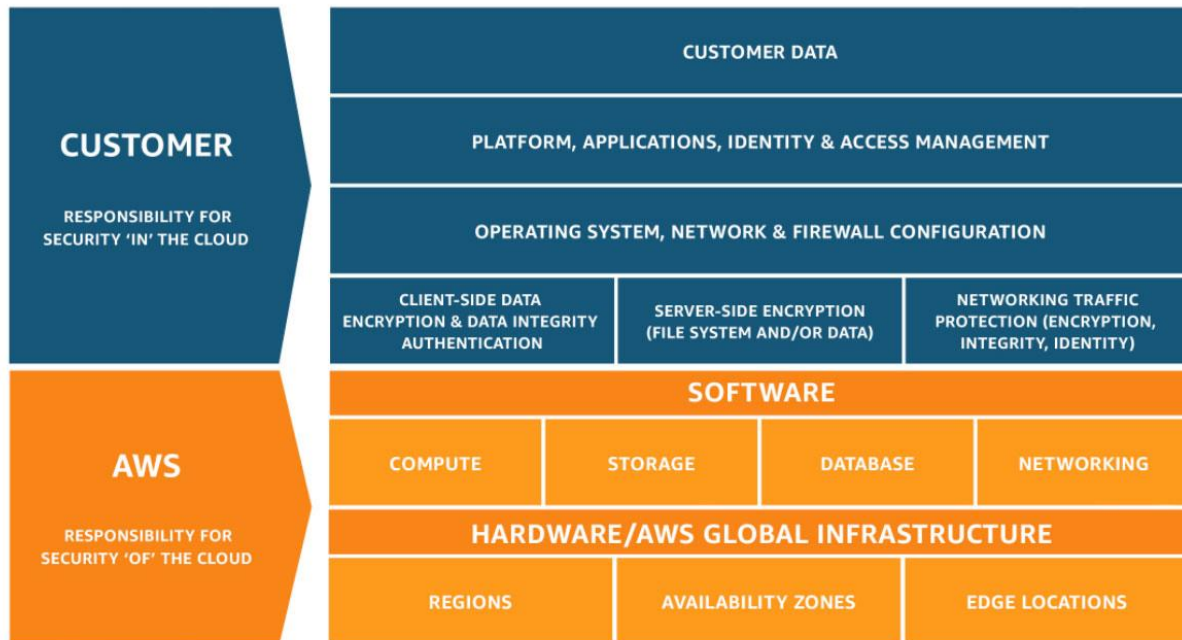


Figure 1 – Shared Responsibility Model

In addition to the direct relationship you have with AWS, there might be other entities that have responsibilities in your particular responsibility model. For example, you might have internal organizational units that take responsibility for some aspects of your operations. You may also have partners or other third parties that develop, manage, or operate some of your cloud technology.

Incident Response in the Cloud

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 consider 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.
- **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 may choose to use Linux `dd` command or a Windows equivalent to make a complete copy of 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 on unknown states.
- **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, plan to fix them. Simulations are safe methods to find gaps and improve processes.

Cloud Security Incidents

Incident Domains

There are three domains within the customer's responsibility where security incidents might occur: service, infrastructure, and application. The difference between the domains is related to the tools you use when you respond. Consider these domains:

- **Service Domain** – Incidents in the service domain affect a customer's AWS account, IAM permissions, resource metadata, billing, and other areas. A service domain event is one that you respond to exclusively with AWS API mechanisms, or have root causes associated with your configuration or resource permissions, and might have related service-oriented logging.
- **Infrastructure Domain** – Incidents in the infrastructure domain include data or network-related activity, such as the traffic to your Amazon EC2 instances within the VPC, processes and data on your Amazon EC2 instances, and other areas. Your response to infrastructure domain events often involves retrieval, restoration, or acquisition of incident-related data for forensics. It likely includes interaction with the operating system of an instance, and in some cases, might also involve AWS API mechanisms.
- **Application Domain** – Incidents in the application domain occur in the application code or in software deployed to the services or infrastructure. This domain should be included in your cloud threat detection and response runbooks, and might incorporate similar responses to those in the infrastructure domain. With appropriate and thoughtful application architecture, you can manage this domain with cloud tools, using automated forensics, recovery and deployment.

In these domains, you must consider the actors who might act against your account, resources, or data. Whether internal or external, use a risk framework to determine what the specific risks are to your organization and prepare accordingly.

In the service domain, you work to achieve your goals exclusively with AWS APIs. For example, handling a data disclosure incident from an Amazon S3 bucket involves API calls to retrieve the bucket's policy, analyzing the S3 access logs, and possibly looking at AWS CloudTrail logs. In this example, your investigation is unlikely to involve data forensic tools or network traffic analysis tools.

In the infrastructure domain, you can use a combination of AWS APIs and familiar digital forensics/incident response (DFIR) software within the operating system of a workstation, such as an Amazon EC2 instance that you've prepared for IR work. Infrastructure domain incidents might involve analyzing network packet captures, disk blocks on an Amazon Elastic Block Store (Amazon EBS) volume, or volatile memory acquired from an instance.

Indicators of Cloud Security Events

There are many security events that you might not classify as incidents, but might still be prudent to investigate. To detect security-related events in your AWS Cloud environment, you can use these mechanisms. Though not an exhaustive list, consider the following examples of some potential indicators:

- **Logs and Monitors** – Review AWS logs, such as Amazon CloudTrail, Amazon S3 access logs, and VPC Flow Logs, and security monitoring services such as [Amazon GuardDuty](#) and [Amazon Macie](#). In addition, use monitors like Amazon Route 53 health checks and [Amazon CloudWatch](#) Alarms. Similarly, use Windows Events, Linux syslog logs, and other application-specific logs that you can generate in your applications.
- **Billing Activity** – A sudden change in billing activity can indicate a security event.
- **Threat Intelligence** – If you subscribe to a third-party threat intelligence feed, you can correlate that information with other logging and monitoring tools to identify potential indicators of events.
- **Partner Tools** – Partners in the AWS Partner Network (APN) offer hundreds of industry-leading products that can help you meet your security objectives. For more information, see [Security Partner Solutions](#) and [Security Solutions in the AWS Marketplace](#).
- **AWS Outreach** – [AWS Support](#) might contact you if we identify abusive or malicious activity. For more information, see the [AWS Response to Abuse and Compromise](#) section.

- **One-Time Contact** – Because it could be your customers, your developers, or other staff in your organization who notice something unusual, it is important to have a well-known, well-publicized method of contacting your security team. Popular choices include ticketing systems, contact email addresses, and web forms. If your organization works with the general public, you might also need to have a public-facing security contact mechanism.

Understanding Cloud Capabilities

AWS offers a wide range of security capabilities that you can use to investigate security events across the domains. For example, AWS provides a number of logging mechanisms, such as AWS CloudTrail logs, Amazon CloudWatch Logs, Amazon S3 access logs, and more. You should consider the services that you're using, and make sure you have enabled the logs that pertain to those services. AWS also offers a [Centralized Logging Solution](#) on AWS Solutions, which can help you to understand how to centralize and store the common types of cloud logs. After you have enabled these logging sources, you must decide how you want to analyze them, such as using [Amazon Athena](#) to query logs held in your Amazon S3 buckets.

Additionally, there are a number of AWS partner products that can simplify your process to analyze these logs, such as those described in the [APN Security Competency program](#). There are also several AWS services that can help you get valuable insights into this data, such as [Amazon GuardDuty](#) (a threat detection service) and [AWS Security Hub](#), which can give you a comprehensive view of your high-priority security alerts and compliance status across AWS accounts. For more information about additional cloud capabilities that you can leverage during your investigations, see [Appendix B: Cloud Capability Definitions](#).

Data Privacy

We know customers care deeply about privacy and data security, and so we implement responsible and sophisticated technical and physical controls designed to prevent unauthorized access to or disclosure of customer content. Maintaining customer trust is an ongoing commitment. You can learn more about AWS data privacy commitments on our [Data Privacy FAQ](#) page.

These intentional, self-imposed controls, limit the ability of AWS to assist in responding within a customer's environment. Because of this, focusing on understanding and building capabilities within the Shared Responsibility Model is key to success in the AWS Cloud. Although enabling logging and monitoring capabilities in your AWS

accounts before an incident occurs is important, there are additional aspects to incident response that are imperative to a successful program.

AWS Response to Abuse and Compromise

Abuse activities are externally observed behaviors of AWS customers' instances or other resources that are malicious, offensive, illegal, or could harm other internet sites. AWS works with you to detect and address suspicious and malicious activities from your AWS resources. Unexpected or suspicious behaviors from your resources can indicate that your AWS resources have been compromised, which signals potential risks to your business.

AWS detects abuse activities in your resources using mechanisms, such as:

- AWS internal event monitoring
- External security intelligence against AWS network space
- Internet abuse complaints against AWS resources

Although the AWS abuse response team aggressively monitors and shuts down unauthorized activity running on AWS, the majority of abuse complaints refer to customers who have legitimate business on AWS. Some examples of common causes of unintentional abuse activities include:

- **Compromised resource** – An unpatched Amazon EC2 instance could be infected and become a botnet agent.
- **Unintentional abuse** – An overly aggressive web crawler might be classified as a denial-of-service attack by some internet sites.
- **Secondary abuse** – An end user of the service provided by an AWS customer might post malware files on a public Amazon S3 bucket.
- **False complaints** – Sometimes internet users mistakenly report legitimate activities as abuse.

AWS is committed to working with AWS customers to prevent, detect, and mitigate abuse, and to defend against future recurrences. We encourage you to review the [AWS Acceptable Use Policy](#), which describes prohibited uses of the web services offered by Amazon Web Services and its affiliates. To support timely response to abuse notifications from AWS, make sure that your AWS account contact information is accurate. When you receive an AWS abuse warning, your security and operational staff should immediately investigate the matter. Delay can prolong the reputation impact and

legal implications to others and to yourself. More importantly, the implicated abuse resource may be compromised by malicious users, and ignoring the compromise could magnify damages to your business.

Prepare – People

Automated processes enable organizations to spend more time focusing on measures to increase the security of their cloud environment and applications. 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, analysts and responders within a security organization still have much to do.

Define Roles and Responsibilities

The skills and mechanisms of incident response are most important when handling new or large-scale events. Because we cannot predict or codify all potential events, humans must do whatever the automation cannot. Handling unclear security events requires cross-organizational discipline, bias for decisive action, and the ability to deliver results. Within your organizational structure, there are probably many people who are responsible, accountable, consulted, or kept informed during an incident. Consider these roles and responsibilities, and whether any third parties must be involved.

Trusted partners might be involved in the investigation or response that provide additional expertise and valuable scrutiny. You might need to contact owners of impacted applications or resources, because they are subject matter experts (SMEs) that can provide information and context. Application owners or SMEs might be required to act in situations where the environment is unfamiliar, has unanticipated complexity, or where the responders do not have access. Application SMEs should practice and become comfortable working with the IR team.

Provide Training

To reduce dependency and decrease response time, make sure that your security teams and responders are educated about cloud services and have opportunities to practice hands-on with the specific cloud platforms that your organization uses. AWS provides a number of training options and learning paths through digital training, classroom training, partners, and certifications. To learn more, see [AWS Training & Certification](#).

Beyond general cloud experience, most customers need to significantly invest in their 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.

Define Response Mechanisms

When you consider your approach to incident response in the cloud, in unison with other teams (such as your legal counsel, leadership, business stakeholders, and others), you must understand what you have and what you need. First, identify stakeholders and relevant contacts, and make sure you have appropriate access to perform the necessary response.

Although the cloud can provide you with greater visibility and capabilities through service APIs, you must document where these are and how to use them. Identify your team's AWS account numbers, the IP ranges of your Virtual Private Clouds (VPCs), corresponding network diagrams, logs, data locations, and data classifications. Many of these technological processes are included in the

Prepare – Technology section. Then, begin documenting your incident response procedures, often referred to as procedures or runbooks, that define the steps to investigate and remediate an incident.

Create a Receptive and Adaptive Security Culture

We have learned that our customers' most successful security teams are cooperative enablers for their business and its developers, who foster a culture that makes sure all stakeholders cooperate and escalate to maintain an agile, high-responsive security posture. Although improving your organization's security culture is not the subject of this paper, you can get relevant intelligence from your non-security staff if they see the security team is receptive. When your security team is open and accessible, with support from leadership, they are more likely to get additional, timely notifications and responses to security events.

In some organizations, staff might fear retribution if they report a security problem. Sometimes they simply don't know how to report an issue.

In other cases, they might not want to waste time, or might be embarrassed to report something as a security incident that is later discovered to not be a problem. From the leadership team down, it is important to promote a culture of acceptance and to invite everyone to be a part of the organization's security. Provide a clear channel for anyone to open a high-severity ticket, whenever they believe there could be a potential risk or threat. Welcome these notifications with an eager and open mind, but more importantly, make it clear to non-security staff that you welcome these notifications. Emphasize that you would rather be over-notified of potential issues, than to receive no notifications at all.

These notifications offer valuable opportunities to practice responsive investigations under stress. They can serve as an important feedback loop while you develop your response procedures.

Predicting Response

Because it is impossible to predict all potential events, you must continue to rely on human analysis. Taking the time to carefully train your staff and prepare your organization helps you to anticipate the unexpected, however, your organization does not have to prepare in isolation. Collaborating with trusted security partners to identify unexpected security events gives organizations the benefit of additional visibility and insight.

Partners and the Window of Response

The journey to the cloud is unique for every organization. However, there are patterns and practices that other organizations have already encountered that a trusted security partner can bring to your attention. 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.

In 1955, Joseph Luft and Harrington Ingham created the Johari Window, an exercise for mapping traits to categories. The window is depicted as a grid that consists of four quadrants, similar to the following diagram.

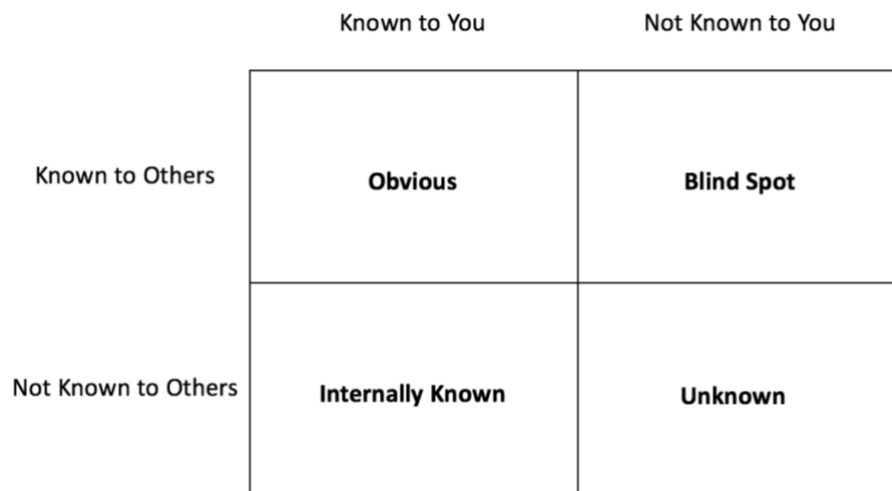


Figure 2 – Window of Response

Although Luft and Ingham's window was not intended for information security, we can adjust the concept to use it as a simple mental model to consider the difficulty in assessing an organization's threats. In our modified concept, the four quadrants are:

- **Obvious** – Risk of which both your team and your partner are aware.
- **Internally Known** – Risk that your team is familiar with, but that your partner isn't. This could mean that you have internal expertise or tribal knowledge.
- **Blind Spot** – Risk that your partner is familiar with, but that your team is not.
- **Unknown** – Risk which neither you or your partner are familiar with.

Although this diagram is simple, it represents the value that having trusted partners can achieve. Most critically, there might be *blind spot* items that you are unaware of, but a partner with the right expertise can bring to your attention. Although you might both be

familiar with those risks in the *obvious* quadrant, your partner could recommend controls and solutions that you are unfamiliar with. Additionally, although you might bring those risks in the *internally known* quadrant to your partner's attention, your partner might also be able to identify optimized controls for mitigating that risk. As you measure yourself for improvement, contact your partner to provide expert advice.

The Unknown

If you've been focused on tailoring alerts, improving your incident response procedures with automation, and improving your security defenses, you might be wondering what to improve next. You might be curious about your unknown risk, as represented in the unknown category of [Figure 2 – Window of Response](#). You can reduce unknown risk through the following methods:

- **Define security assertions** – What are some truths that you can assert? What are the security primitives that should absolutely be true in your environment? Clearly defining these allows you to search for the inverse. This is something that's easier to do early in your cloud journey, rather than attempting to reverse engineer your security assertions later.
- **Education, communication, and research** – Create cloud security experts on your staff or include expert partners to help scrutinize your environment. Challenge your assumptions, and be wary of subtle reasoning. Create feedback loops in your processes and offer mechanisms for your engineering teams to communicate with security teams. You can also expand your approach to monitor relevant security mailing lists and information security disclosures.
- **Reducing attack surface** – Improve your defense to avoid risk and give yourself more time against unknown attacks. Block and slow down attackers, and force them to be noisy.
- **Threat intelligence** – Subscribe to a continuous feed of current and relevant threats, risk, and indicators from around the world.
- **Alerts** – Generate notifications that alert you to unusual, malicious, or expensive activities. For example, you might create a notification for activities that occur in regions or services that you do not use.
- **Machine learning** – Leverage machine learning to identify complex anomalies for a specific organization or individual personas. To help you identify unusual behaviors, you can also profile the normal characteristics of your networks, users, and systems.

Amazon Web Services provides you with the tools and services to manage these issues yourself. Using machine learning to identify malicious patterns is a well-researched field of study, with patterns that are implemented by customers, AWS Professional Services, AWS Partners, and through AWS services such as Amazon GuardDuty and Amazon Macie. Some of these patterns have been discussed at AWS re:Invent conference sessions. For more information, see the [Media](#) section.

Customers are also expanding their traditionally business-centric data lakes to leverage similar architecture patterns when they develop security data lakes. Security operations teams are also expanding their use of traditional logging and monitoring tools, such as Amazon Elasticsearch Service and Kibana, to big data architectures.

Those customers are collecting internal data from AWS CloudTrail event logs, VPC Flow Logs, Amazon CloudFront access logs, database logs, and application logs, and then combine this data with public data and threat intelligence. With this valuable data, customers have expanded to include data science and data engineering skills on their security operations teams to leverage tools such as Amazon EMR, Amazon Kinesis Data Analytics, Amazon Redshift, Amazon QuickSight, AWS Glue, Amazon SageMaker, and Apache MXNet on AWS to build custom solutions that identify and predict anomalies that are unique to their business.

Prepare – Technology

Prepare Access to AWS Accounts

During an incident, your incident response teams must have access to the environments and resources involved in the incident. Make sure that your teams have appropriate access to perform their duties before an event occurs. To do that, 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. Your team members' authentication and authorization should be documented and tested well before an event occurs to make sure they can perform a timely response without delays.

At this stage, you must work closely with your cloud architecture and development teams to understand what level of access is necessary for responders. Identify and discuss the AWS account strategy and cloud identity strategy with your organization's cloud architects to understand what authentication and authorization methods are configured, for example:

- Federation – A user assumes an IAM role in an AWS account from an Identity Provider.
- Cross-account access – A user assumes an IAM role between multiple AWS accounts.
- Authentication – A user authenticates as an AWS IAM user created within a single AWS account.

These options define the technical choices for authentication to AWS, and how you may gain access during a response, but some organizations might rely on another team or a partner to assist in the response. User accounts that are 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, and they should not be used for daily activities.

Before you create new access mechanisms, work with your cloud teams to understand how your AWS accounts are organized and governed. Many customers use AWS Organizations to help centrally manage billing, share resources across their AWS accounts, and control access, compliance, and security. A core feature of Organizations is that it can be leveraged to apply [Service Control Policies](#) to groups of accounts, which enables you to gain policy management at scale. For additional information about

implementing governance mechanisms at scale, see [AWS Governance at Scale](#). After you understand how your organization has organized and governed your AWS accounts, consider the following generalized response patterns to assist in identifying which approaches are right for your organization.

Indirect Access

If you use *indirect access*, your account owners or application teams are required to perform authorized remediations in their AWS accounts with tactical guidance from the incident response team who are your security experts. This method is a slower and more complex way to execute tasks, but it can be successful when the responders are unfamiliar with the account or cloud environment.

Direct Access

To give incident responders *direct access*, you deploy an AWS IAM role into the AWS accounts that your security engineers or incident responders can assume during a security event. The incident responder authenticates either through a normal federated process, or through a special emergency process, if the incident impacts your normal authentication process. The permissions you give the incident response IAM role depend on the actions you anticipate the responders to perform.

Alternative Access

If you believe a security event is impacting your security, identity, or communication systems, you might need to seek alternative mechanisms and access to investigate and remediate the impact. By using a new, purpose-built AWS account, your responders can collaborate and work from an alternate, secure infrastructure.

For example, responders can leverage new infrastructure launched in the cloud, such as remote workstations using [Amazon WorkSpaces](#) and email services provided by [Amazon WorkMail](#). You must prepare appropriate access controls (using IAM policies) to delegate access so that your secure, alternative AWS account can assume permissions for the impacted AWS account.

After you have delegated appropriate access, you can use the AWS APIs in the affected account to share relevant data, such as logs and volume snapshots, to perform investigative work in the isolated environment. For more information about this cross-account access, see [Tutorial: Delegate Access Across AWS Accounts Using IAM Roles](#).

Automation Access

As you migrate to using automation to respond to security events, you must create IAM roles specifically for your automation resources to use (such as Amazon EC2 instances or AWS Lambda functions). These resources can then assume the IAM roles and inherit the permissions assigned to the role. Instead of creating and distributing AWS credentials, you delegate permission to your AWS Lambda function or Amazon EC2 instance. The AWS resource automatically receives a set of temporary credentials and uses them to sign API requests.

You can also consider a secure method for your automation or tooling to authenticate and execute within the operating system of your Amazon EC2 instance. Though there are many tools that can perform this automation, consider 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.

The AWS Systems Manager Agent (SSM Agent) is installed by default on some Amazon EC2 Amazon Machine Images (AMIs), such as for Windows Server and Amazon Linux. However, you might need to manually install the agent on other versions of Linux and hybrid instances. Whether you use Run Command or another tool, complete any prerequisite setup and configuration before you receive your first security-related alert to investigate.

Managed Services Access

Your organization may already be partnered with an information technology provider that manages services and solutions on your behalf. These partners have a shared responsibility in supporting the security of your organization, and it's important to clearly understand this relationship before an anomaly occurs. Whether you already work with an [AWS Managed Service Provider \(MSP\) Partner](#), or AWS Managed Services, or a managed security services partner, you must identify the responsibilities of each partner as they relate to your cloud environments, what access the providers already have to your cloud services, what access they need, and points-of-contact or escalation paths for when you need their assistance. Finally, you should practice this with your partner to make sure that your response plans are predictable and successful.

For example, AWS Managed Services (AMS) provides ongoing management of your AWS infrastructure so that you can focus on your applications. By implementing best practices to maintain your infrastructure, AWS Managed Services helps to reduce your operational overhead and risk. AWS Managed Services automates common activities

such as change requests, monitoring, patch management, security, and backup services, and provides full-lifecycle services to provision, run, and support your infrastructure.

As an Infrastructure Operator, AMS takes responsibility for deploying a suite of security detective controls and providing first line response to them, 24/7 using a follow-the-sun model. When an alert is triggered, AMS follows a standard set of runbooks to ensure a consistent response. These runbooks are shared with AMS customers during onboarding, so they can develop and coordinate response with AMS. AMS encourages the joint execution of security response simulations with customers to develop operational muscle before a real incident occurs.

Prepare Processes

Once appropriate access has been provisioned and tested, your incident response team must define and prepare the related processes necessary for investigation and remediation. This stage is effort intensive, because you must sufficiently plan the appropriate response to security events within your cloud environments.

Work closely with your internal cloud services teams and partners to identify the tasks required to ensure that these processes are possible. Collaborate or assign each other response activity tasks and ensure necessary account configurations are in place. We recommend preparing processes and prerequisite configurations in advance to give your organization the following response capabilities.

Decision Trees

Sometimes, different conditions can require different actions or steps. For example, you might take different actions based on the type of AWS account (development versus production), the tags of the resources, the AWS Config rules compliance status of those resources, or other inputs.

To support you in the creation and documentation of these decisions, we recommend you draft a decision tree with your other teams and stakeholders. Similar to a flow chart, a decision tree is a tool that can be leveraged to support decision making, helping to guide you to determine the optimal actions and outcomes based on potential conditions and inputs, including probabilities.

Use Alternative Accounts

Although responding to an event in the impacted account might be required, it is ideal to investigate data outside of the affected account. Some customers have a process for creating separate, isolated AWS account environments, using templates that preconfigure the resources they must provision. These templates are deployed through a service, such as AWS CloudFormation or Terraform, which provides an easy method to create a collection of related AWS resources and provision them in an orderly and predictable fashion.

Preconfiguring these accounts using templated mechanisms helps to remove human interactions during the initial stages of an incident and ensure the environment and resources are prepared in a repeatable and predictable manner, which can be verified by an audit. In addition, this mechanism also increases the ability to maintain security and containment of data in the forensics environment.

This approach requires you to work with your cloud services and architect teams to determine an appropriate AWS account process that can be used for investigations. For example, your cloud services teams could use [AWS Organizations](#) to generate new accounts and assist you in preconfiguring those accounts using a templated or scripted method.

View or Copy Data

Responders require access to logs or other evidence to analyze and must ensure that they have the ability to view or copy data. At a minimum, the IAM permission policy for the responders should provide read-only access so that they can investigate. To enable appropriate access, you might consider some pre-built AWS Managed Policies, such as [SecurityAudit](#) or [ViewOnlyAccess](#) appropriate access.

For example, responders might want to make a point-in-time copy of data, such as the AWS CloudTrail logs, from an Amazon S3 bucket in one account to an Amazon S3 bucket in another account. The permissions provided by the ReadOnlyAccess managed policy, for example, enable the responder to perform these actions. To understand how to use the AWS Command Line Interface (CLI) to perform this, see [How can I copy objects between Amazon S3 buckets](#).

Sharing Amazon EBS Snapshots

Many customers use Amazon Elastic Block Store (Amazon EBS) snapshots as part of their investigation for security events that involve their Amazon EC2 instances.

Snapshots of Amazon EBS volumes are incremental backups. For more information about Amazon EBS incremental snapshots, see [Amazon EBS Snapshots](#).

To perform an investigation of an Amazon EBS volume in a separate, isolated account, you must modify the permissions of the snapshot to share it with the AWS accounts that you specify. Users that you have authorized can use the snapshots you share as the basis to create their own EBS volumes, while your original snapshot remains unaffected. For more information, see [Sharing an Amazon EBS Snapshot](#).

If your snapshot is encrypted, you must also share the custom KMS Customer Managed Key (CMK) used to encrypt the snapshot. You can apply cross-account permissions to a custom CMK either when it is created or at a later time. Snapshots are constrained to the region in which they were created, but you can share a snapshot with another region by copying the snapshot to that region. For more information, see [Copying an Amazon EBS Snapshot](#).

Sharing Amazon CloudWatch Logs

Logs that are recorded within Amazon CloudWatch Logs, such as Amazon VPC flow logs, can be shared with another account (such as your centralized security account) through a CloudWatch Logs subscription. For example, the log event data can be read from a centralized Amazon Kinesis stream to perform custom processing and analysis. Custom processing is especially useful when you collect logging data from across many accounts. Ideally, create this configuration early in your cloud journey, before a security-related event occurs. For more information, see [Cross-Account Log Data Sharing with Subscriptions](#).

Use Immutable Storage

When copying logs and other evidence to an alternative account, make sure that the replicated data is protected. However, in addition to protecting the secondary evidence, you must protect the integrity of the data at the source. Known as *immutable* storage, these mechanisms protect the integrity of your data by preventing the data from being tampered with or deleted.

Using the native features of Amazon S3, you can configure an Amazon S3 bucket to protect the integrity of your data. By managing access permissions with S3 bucket policies, configuring S3 versioning, and enabling [MFA Delete](#), you can restrict how data

can be written or read. This type of configuration is useful for storing investigation logs and evidence, and is often referred to as *write once, read many* (WORM). You can also protect the data by using server-side encryption with AWS Key Management Service (KMS) and verifying that only appropriate IAM principals are authorized to decrypt the data.

Additionally, if you want to securely keep data in a long-term storage after the investigation is completed, consider moving the data from Amazon S3 to [Amazon S3 Glacier](#) using object lifecycle policies. Amazon S3 Glacier is a secure, durable, and extremely low-cost cloud storage service for data archiving and long-term backup. It is designed to deliver 99.999999999% durability, and provides comprehensive security and compliance capabilities that can help meet even the most stringent regulatory requirements.

Moreover, you can protect the data in Amazon S3 Glacier by using the [Amazon S3 Glacier Vault Lock](#), which allows you to easily deploy and enforce compliance controls for individual Amazon S3 Glacier vaults with a vault lock policy. You can specify security controls, such as WORM, in a vault lock policy and lock the policy from future edits. Once locked, the policy can no longer be changed. Amazon S3 Glacier enforces the controls set in the vault lock policy to help achieve your compliance objectives, such as for data retention. You can deploy a variety of compliance controls in a vault lock policy using the AWS Identity and Access Management (IAM) policy language.

Launch Resources Near the Event

For responders who are new to the cloud, it can be tempting to try to conduct cloud investigations on-premises where your existing tools are located. In our experience, AWS customers who respond to incidents using cloud technologies achieve better results—isolations can be automated, copies can be made more easily, evidence is ready for analysis sooner, and the analysis can be completed faster.

The best practice is to perform investigations and forensics in the cloud, where the data is, rather than attempting to transfer the data to a data center before you investigate. You can use the secure compute and storage capabilities of the cloud practically anywhere in the world to perform the secure response operations. Many customers choose to pre-build a separate AWS account that is ready to perform an investigation, though there might be cases where you choose to operate your analysis in the same AWS account. If your organization is expected to retain records for compliance and legal reasons, it might be prudent to maintain separate accounts for long-term storage and legal activities.

It is also a best practice to perform the investigation in the same AWS Region where the event occurred, rather than replicating the data to another AWS Region. We recommend this practice primarily because of the additional time required to transfer the data between regions. For each AWS Region you choose to operate in, make sure that both your incident response process and the responders abide by the relevant data privacy laws. If you do need to move data between AWS Regions, consider the legal implications of moving data between jurisdictions. It is generally a best practice to keep the data within the same national jurisdiction.

If you believe a security event is impacting your security, identity, or communication systems, you might need to seek alternative mechanisms and access to investigate and remediate the impact. AWS offers you the ability to quickly launch new infrastructure that can be used for secure, alternate work environments. For example, while you investigate the potential severity of the situation, you might want to create a new AWS account with the secure tools for your legal counsel, public relations, and security teams to communicate and continue working. Services such as AWS WorkSpaces (for virtual desktops), AWS WorkMail (for email), and AWS Chime (for communication) can provide your response teams, leadership, and other participants with the capabilities and connectivity they need to communicate, investigate, and remediate an issue.

Isolate Resources

In the course of your investigation, you might need to isolate resources as part of your response to a security anomaly. The intention behind isolating resources is to limit the potential impact, prevent further propagation of affected resources, limit the unintended exposure of data, and prevent further unauthorized access.

As with any response, other business, regulatory, legal, or other considerations can apply. Make sure to weigh your intended actions against expected and unexpected consequences. If your cloud teams use resource tags, these tags can help you identify the criticality of the resource or the owner to contact.

Launch Forensic Workstations

Some of your incident response activities might include analyzing disk images, file systems, RAM dumps, or other artifacts that are involved in an incident. Many customers build a customized forensic workstation that they can use to mount copies of any affected data volumes (known as EBS snapshots). To do so, follow these basic steps:

1. Choose a base Amazon Machine Image (AMI) (such as Linux or Windows) that can be used as a forensic workstation.
2. Launch an Amazon EC2 instance from that base AMI.
3. Harden the operating system, remove unnecessary software packages, and configure relevant auditing and logging mechanisms.
4. Install your preferred suite of open source or private toolkits, as well as any vendor software and packages you need.
5. Stop the Amazon EC2 instance and create a new AMI from the stopped instance.
6. Create a weekly or monthly process to update and rebuild the AMI with the latest software patches.

After the forensic system is provisioned using an AMI, your incident response team can use this instance to create a new AMI to launch a new forensic workstation for each investigation. The process for launching the AMI as an Amazon EC2 instance can be preconfigured to simplify the deployment process. For example, you can create a template of the forensic infrastructure resources you need within a text file and deploy it into your AWS account through AWS CloudFormation.

When your resources are available to be deployed quickly from a template, your well-trained forensic experts are able to use new forensic workstations for each investigation, instead of reusing infrastructure. With this process, you can make sure that there is no cross-contamination from other forensic examinations.

Instance Types for Workstations

Amazon EC2 provides a wide selection of instance types that are optimized for different use cases. Instance types comprise varying combinations of CPU, memory, storage, and networking capacity, and give you the flexibility to choose the appropriate mix of resources for your applications. Each instance type includes one or more instance sizes, which enables you to scale your resources to the requirements of your target workload.

AWS enhanced networking uses single root I/O virtualization (SR-IOV) to provide high-performance networking capabilities on [supported instance types](#). SR-IOV is a method of device virtualization that provides higher I/O performance and lower CPU utilization when compared to traditional virtualized network interfaces. Enhanced networking provides higher bandwidth, higher packet per second (PPS) performance, and consistently lower inter-instance latencies. There is no additional charge for using enhanced networking. For information about which instance types support 10 or 25 Gbps network speeds, and other advanced capabilities, see [Amazon EC2 Instance Types](#).

Cloud Provider Support

AWS Support

[AWS Support](#) offers a range of plans that provide access to tools and expertise that support the success and operational health of your AWS solutions. All support plans provide 24x7 access to customer service, AWS documentation, whitepapers, and support forums. If you need technical support and more resources to help plan, deploy, and optimize your AWS environment, you can select a support plan that best aligns with your AWS use case.

You should consider the [Support Center](#) in the AWS Console as the central point of contact to get support for issues that affect your AWS resources. Access to AWS Support is controlled by AWS Identity and Access Management (IAM). For more information about getting access to AWS support features, see [Accessing Support](#).

If you need to report abuse of Amazon EC2, contact the AWS Abuse team at: <https://aws.amazon.com/forms/report-abuse>

DDoS Response Support

A Denial of Service (DoS) attack makes your website or application unavailable to end users. Attackers use a variety of techniques that consume network or other resources, disrupting access for legitimate end users. In its simplest form, a DoS attack against a target is executed by a lone attacker from a single source.

In a Distributed Denial of Service (DDoS) attack, an attacker uses multiple sources, which may be compromised or controlled by a group of collaborators, to orchestrate an attack against a target. In a DDoS attack, each of the collaborators or compromised hosts participates in the attack, generating a flood of packets or requests to overwhelm the intended target.

AWS offers customers [AWS Shield](#), which provides a managed DDoS protection service that safeguards web applications running on AWS. AWS Shield provides always-on detection and automatic inline mitigations that minimize application downtime and latency, so there is no need to engage AWS Support to benefit from DDoS protection. There are two tiers of AWS Shield: Standard and Advanced.

All AWS customers benefit from the automatic protections of AWS Shield Standard, at no additional charge. AWS Shield Standard defends against most common, frequently occurring network and transport layer DDoS attacks that target your website or applications. When you use AWS Shield Standard with Amazon CloudFront and Amazon Route 53, you receive comprehensive availability protection against all known infrastructure (Layer 3 and 4) attacks.

For higher levels of protection against attacks targeting your web applications running on [Amazon Elastic Compute Cloud \(Amazon EC2\)](#), [Elastic Load Balancing \(ELB\)](#), [Amazon CloudFront](#), and [Amazon Route 53](#) resources, you can subscribe to AWS Shield Advanced. Additionally, AWS Shield Advanced gives you 24x7 access to the AWS DDoS Response Team (DRT). For more information about AWS Shield Standard and AWS Shield Advanced, see [AWS Shield](#).

Simulate

Security Incident Response Simulations

Security Incident Response Simulations (SIRS) are internal events that provide a structured opportunity to practice your incident response plan and procedures during a realistic scenario. SIRS events are fundamentally about being prepared and iteratively improving your response capabilities. Some of the reasons customers find value in performing SIRS 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 with laser focus
- 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.

Simulation Steps

Regardless of whether you design your own SIRS, or have a trusted partner to provide the groundwork, simulations generally follow these steps:

1. **Find an issue of importance** – Define the trigger that should cause a response.
2. **Identify skilled security engineers** – A simulation requires a builder and a tester.
3. **Build a realistic model system** – The simulation must be realistic and appropriate. If it is not realistic, participants might not value the exercise. If it is too minimal, the exercise might be deemed trivial. Start with simple exercises and work towards a full event.
4. **Build and test the scenario elements** – Relevant simulation material might need to be built, such as logging artifacts, email notifications and alerts, and potential runbooks.
5. **Invite other security individuals and cross-organizational participants** – Invite everyone who needs to train and participate. If your general legal counsel, executives, and public relations have a part in the simulation, you should also invite them.
6. **Run the simulation**– Choose if your staff should expect the SIRS event, or if the simulation should remain unannounced.
7. **Celebrate, measure, improve, and repeat** – The simulation has factors of stress, and so it is important to encourage and celebrate the efforts of your participants. After encouragement comes the opportunity to measure, improve, and iterate for the next simulation. AWS encourages you to make a habit of these activities.

Important – If you are planning to perform a Security Incident Response Simulation, go to the [Penetration Testing page](#) and review the *Other Simulated Events* section for the latest information on how to proceed.

Simulation Examples

Security simulations must be realistic to provide the expected value. When you or your partners work toward creating your own simulations, always consider past, real-world events as a valuable source for potential simulation exercises. Here are a few examples that AWS customers have found useful to use for their initial simulations:

- Unauthorized changes to network configuration or resources
- Credentials that were mistakenly exposed publicly due to developer misconfiguration
- Sensitive content that was mistakenly made publicly-accessible by developer misconfiguration
- Isolation of a web server that is communicating with suspected malicious IP addresses

In addition to the valuable, experiential learning, performing SIRS activities generates outputs, such as lessons learned, that you can use as inputs into the next process of your program—iteration.

Iterate

The previous section defined some of the benefits of SIRS activities. Among these advantages was gaining agility through incremental improvements. Simulations should generate valuable outcomes that you can leverage to improve your security response. They provide a feedback loop to the organization, about what is working and what is not working. With this knowledge, you can incrementally create new procedures or update existing ones to improve your response.

Runbooks

When a security anomaly is detected, containing the event and returning to a known good state are important elements of a response plan. As an example, if the anomaly occurred because of a security misconfiguration, the remediation might be as simple as removing the variance through a redeployment of the resources with the proper configuration. To do this, you will need to plan ahead and define your own security response procedures, which are often called runbooks.

A *runbook* is the documented form of an organization's procedures for conducting a task or series of tasks. This documentation is usually stored either in an internal digital system or on printed paper. You might currently have incident response runbooks, or you might need to create them to be compliant to a security assurance framework. However, when you manually follow written runbooks, you increase the potential that you will make mistakes. Instead, we recommend that you automate all of your repeatable tasks. Automation frees your response team from common tasks, and makes them available for more important tasks, such as correlating events, practicing in simulations, devising new response procedures, performing research, developing new skills, and testing or building new tools. However, before you can decompose the tasks into programmable logic and iterate towards proper automation, you must start by writing a runbook.

Creating Runbooks

To create runbooks for the cloud, we recommend that you first focus on the alerts you currently generate. If you generate an alert, it is important to investigate it. Start by defining 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.

After you determine the best solution, 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 speeds 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 them being executed by a human responder.

Getting Started

If you're not sure where to start, consider beginning with the alerts that could be generated by [AWS Trusted Advisor](#) and [AWS Config Rules](#). Then, focus on events generated by [Amazon GuardDuty](#) and [Amazon Macie](#). Information about Amazon GuardDuty findings are available on the [Amazon GuardDuty User Guide](#). Information about Amazon Macie alerts are available on the Amazon Macie service pages in the AWS Management Console.

Both Amazon GuardDuty and Amazon Macie send notifications through Amazon CloudWatch Events when any change in the findings or alerts occurs. This includes newly generated alerts and updates to existing alerts. You can set up the Amazon CloudWatch Events rules to trigger AWS Lambda functions to perform event-driven response. For more information, see the [Event-Driven Response](#) section.

Automation

Automation is a force multiplier, which means it scales the efforts of your responders to match the speed of the organization. Moving from manual processes to automated processes enables you to spend more time increasing the security of your AWS Cloud environment.

Automating Incident Response

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 the 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.

However, your objective should be to further reduce the time gap between detective mechanisms and responsive mechanisms. Historically, this time gap can take hours, days, or even months. An [Incident Response survey by SANS in 2016](#) found that 21% of respondents stated their time to detection took two to seven days, and only 29% of respondents were able to remediate events within the same time frame. In the cloud, you can reduce that response time gap to seconds by building event-driven response capabilities.

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), the response procedure is to enable the service again and investigate the user that disabled the AWS CloudTrail logging. Instead of performing these steps manually in the AWS Management Console, you can programmatically enable the logging again (through the `cloudtrail:StartLogging` API). If you implement this with code, your response objective is to perform this task as quickly as possible and notify the responders that the response was performed.

You can decompose the logic into simple code to run in an AWS Lambda function to perform these tasks. You can then use Amazon CloudWatch Events to monitor for the specific `cloudtrail:StopLogging` event, and invoke the function if it occurs. When this AWS Lambda responder function is invoked by Amazon CloudWatch Events, you can pass it the details of the specific event with the information of the principal that disabled AWS CloudTrail, when it was disabled, the specific resource that was affected, and other relevant information. You can use this information to essentially perform a log dive, and then generate a notification or alert with only the specific values that a response analyst would require.

Ideally, the goal of event-driven response is for the Lambda responder function to perform the response tasks and then notify the responder that the anomaly has been successfully resolved with any pertinent contextual information. It is then up to the human responder to decide how to determine why it occurred and how future reoccurrences might be prevented. This feedback loop drives further security improvement into your cloud environments. To achieve this objective, you must have a culture that enables your security team to work closer with your development and operations teams.

Incident Response Examples

Service Domain Incidents

You can handle many service domain incidents exclusively through AWS APIs.

Identities

Amazon Web Services provides APIs to cloud services that are used by millions of customers to build new applications and drive business outcomes. These APIs can be invoked through many methods, such as by software development kits (SDKs), the AWS CLI, and the AWS Management Console. To interact with AWS through these methods, the AWS Identity and Access Management (IAM) service helps you securely control access to AWS resources. You can use IAM to control who is authenticated (signed in) and authorized (has permissions) to use resources at the Account Level. For a list of AWS services that you can use with IAM, see [AWS Services That Work with IAM](#).

When you first create an AWS account, you begin with a single sign-on (SSO) identity that has complete access to all AWS services and resources in the account. This identity is called the AWS account root user and is accessed by signing in with the email address and password that you used to create the account. We strongly recommend that you do not use the root user for your everyday tasks, and particularly not for administrative tasks. Instead, we recommend that you follow the best practice of using the root user only to create your first IAM user, securely store the root user credentials, and use them to perform only a few account and service management tasks. For more information, see [Create Individual IAM Users](#).

Although these APIs provide value to millions of customers, some of them can be abused if the wrong individuals get access to your IAM account or root credentials. For example, you can use the APIs to enable logging within your account, such as AWS CloudTrail. However, if attackers get your credentials, they can also use the API to disable these logs. You can prevent this type of abuse by configuring appropriate IAM permissions that follow a least privilege model, and by properly protecting your IAM credentials. For more information, see [IAM Best Practices](#) in the *AWS Identity and Access Management User Guide*. If this type of event does occur, there are multiple detective controls to identify that your AWS CloudTrail logging was disabled, including AWS CloudTrail, AWS Config, AWS Trusted Advisor, and AWS CloudWatch Events.

Resources

Other features that can be abused or misconfigured vary from organization to organization based on how each customer operates in the cloud. For example, some organizations intend to make certain data or applications publicly accessible, while others keep their applications and data internal and confidential. Not all security events are malicious in nature; some events might result from unintentional or improper configurations. Consider which APIs or features have a high impact to your organization, and whether you use them frequently or infrequently.

You can identify many security misconfigurations using tools and services. For example, AWS Trusted Advisor provides a number of checks for best practices. Partners in the AWS Partner Network (APN) also offer hundreds of industry-leading products that are equivalent, identical to, or integrate with existing controls in your on-premises environments. A number of these products and solutions have been prequalified by the [AWS Partner Competency Program](#). We encourage you to visit the [Configuration and Vulnerability Analysis](#) section of the APN Security Competency program to browse these solutions and to determine if they can satisfy your requirements.

Infrastructure Domain Incidents

The infrastructure domain typically includes your application's data or network-related activity, such as the traffic to your Amazon EC2 instances within the VPC and the processes running in your Amazon EC2 instance operating systems.

For example, assume that your monitoring solution notified you of a potential security anomaly on your Amazon EC2 instance. The following are common steps to address this issue:

1. **Capture** the metadata from the Amazon EC2 instance, before you make any changes to your environment.
2. **Protect** the Amazon EC2 instance from accidental termination by [enabling termination protection for the instance](#).
3. **Isolate** the Amazon EC2 instance by switching the VPC Security Group or explicitly denying network traffic to the IP address of the instance with the Network Access Control List.
4. **Detach** the Amazon EC2 instance from any [AWS Auto Scaling](#) groups.
5. **Deregister** the Amazon EC2 instance from any related [Elastic Load Balancing](#) service.

6. **Snapshot** the Amazon EBS data volumes that are attached to the EC2 instance for preservation and follow-up investigations.
7. **Tag** the Amazon EC2 instance as quarantined for investigation, and add any pertinent metadata, such as the trouble ticket associated with the investigation.

You can perform all of the preceding steps using the AWS APIs, AWS SDKs, the AWS CLI, and the AWS Management Console. To interact with AWS using these methods, the AWS Identity and Access Management (IAM) service helps you securely control access to AWS resources. You use IAM to control who is authenticated and authorized to use resources at the Account Level. The IAM service provides the authentication and authorization for you to perform these actions and interact with the service domain.

A snapshot of an Amazon EBS volume is a point-in-time, block-level copy of an EBS data volume, which occurs asynchronously and might take time to complete. You can create new EBS volumes from these copies and mount them to the forensic EC2 instance for deep analysis offline by forensic investigators. The following diagram shows a simplified version of the outcome, and does not describe all of the network components (such as subnets, routing tables, and network access control lists).

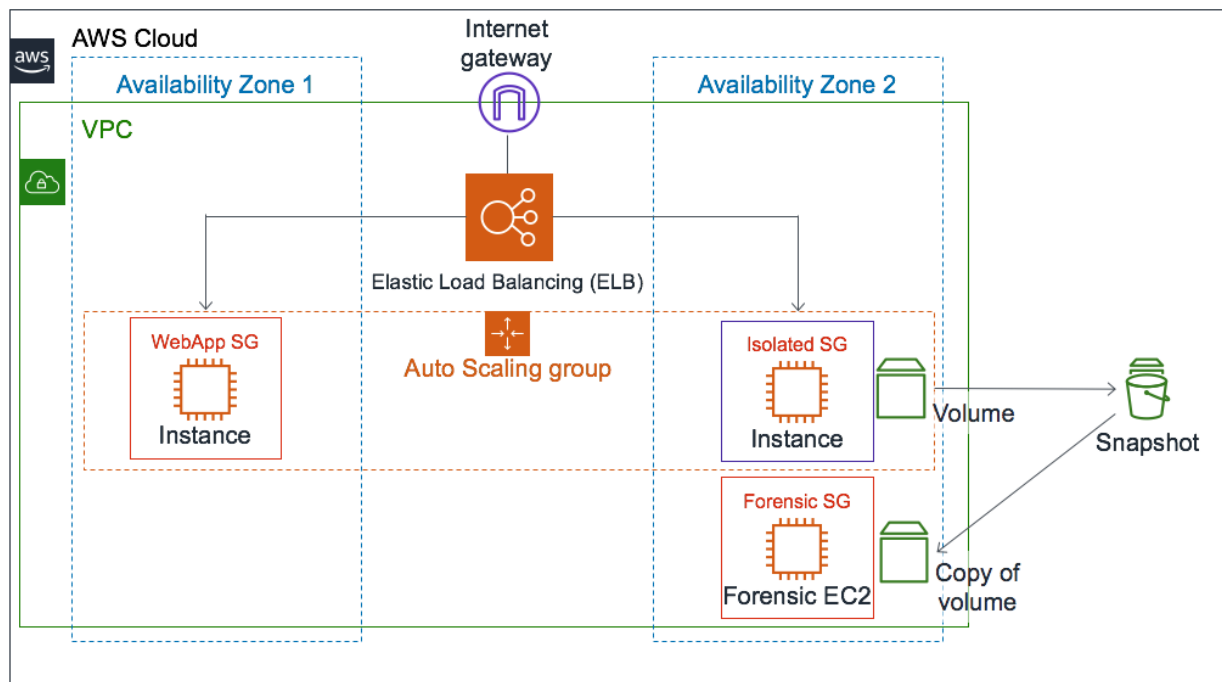


Figure 3 – EC2 Instance Isolation and Snapshots

Investigation Decisions

At this point, you can choose between an offline investigation (immediately shut down the instance) or an online investigation (keep the instance running). One advantage to the offline investigation is that after the instance is shut down, it can no longer affect the existing environment. Additionally, you can create a copy of the affected instance from the EBS snapshots, and review it in an isolated AWS account with an isolated environment that is designed specifically for your investigation. However, you can choose to not shut down the instance immediately, if an *online* investigation enables you to potentially capture volatile evidence from the host operating system, such as memory or network traffic.

Capturing Volatile Data

Although you might not choose to perform the online investigation, it is important to understand the necessary mechanisms to capture volatile data from an instance. An online investigation requires interaction with your operating system that is running on the Amazon EC2 instance. In this scenario, you need more than the AWS IAM service to execute tasks on an Amazon EC2 instance. Although you could authenticate directly to the machine using a standard method—such as Linux secure shell (SSH) or Windows remote desktop (RDP)—manual interaction with the operating system is not a best practice. We recommend that you programmatically use an automation tool to execute tasks on a host.

Using AWS Systems Manager

The [AWS Systems Manager Run Command](#) helps you to remotely and securely perform on-demand changes running Linux shell scripts and Windows PowerShell commands on a targeted instance. Although you can invoke Run Command through permissions in the AWS IAM service, you must first activate your Amazon EC2 instances as managed instances, install the SSM Agent on your machines (if not installed by default), and configure the AWS IAM permissions. If you are interested in using Run Command for automation or response activities, make sure to complete the prerequisite activities before you have to perform an investigation.

Systems Manager, which includes Run Command, is integrated with AWS CloudTrail, a service that captures API calls made by or on behalf of Systems Manager and delivers the log files to an Amazon S3 bucket that you specify. Using the information collected by AWS CloudTrail, you can determine what request was made, the source IP address that made the request, who made the request, when it was made, and more. CloudTrail

creates logs of all Systems Manager API actions, including API requests to execute commands using Run Command or to create Systems Manager documents.

You can use the AWS Systems Manager Run Command service to invoke the SSM Agent that executes Linux shell scripts and Windows PowerShell commands. These scripts can load and execute specific tools to capture additional data from the host, such as the Linux Memory Extractor (LiME) kernel module. You can then transfer the memory capture to your forensic Amazon EC2 instance in the VPC network, or to an Amazon S3 bucket for durable storage.

Automating the Capture

One method to invoke the SSM Agent is to target the Run Command through Amazon CloudWatch Events when the instance is tagged with a specific tag. For example, if you apply the `Response=Isolate+MemoryCapture` tag to an affected instance, you can configure Amazon CloudWatch Events to trigger two actions: 1) a Lambda function that performs the isolation activities, and 2) a Run Command that executes a shell command to export the Linux memory through the SSM Agent. This tag-driven response is another method of event-driven response.

Conclusion

As you continue your cloud journey, it is important for you to consider the aforementioned fundamental security incident response concepts for your AWS environment. You can combine the available controls, cloud capabilities, and remediation options, to help you improve the security of your cloud environment. You can start small and iterate as you adopt automation capabilities that improve your response speed, so you are better prepared when security events occur.

Contributors

Contributors to this document include:

- Joshua Du Lac, Sr. Solutions Architect, Security Specialist, AWS Solutions Architecture
- Nathan Case, Sr. Solutions Architect, Security Specialist, AWS Solutions Architecture
- Paco Hope, Principal Consultant, AWS ProServe

Additional Resources

For additional information, see:

- [AWS Well-Architected](#)
- [Security Perspective of the AWS Cloud Adoption Framework \(CAF\)](#)
- [AWS Centralized Logging Solution](#)
- [Visualize AWS CloudTrail Logs using AWS Glue and Amazon QuickSight](#)
- [How to Visualize and Refine Your Network's Security by Adding Security Group IDs to Your VPC Flow Logs](#)
- [How to Monitor Host-Based Intrusion Detection System Alerts on Amazon EC2 Instances](#)
- [Store and Monitor OS & Application Log Files with Amazon CloudWatch](#)
- [Identity and Access Management in Amazon S3](#)
- [Using Versioning \(Amazon S3\)](#)
- [Using MFA Delete](#)
- [Protecting Data Using Server-Side Encryption with AWS KMS–Managed Keys \(SSE-KMS\)](#)
- [Incident Response with AWS Console and CLI](#)

Media

- [AWS re:Invent 2014 \(SEC402\): Intrusion Detection in the Cloud](#)
- [AWS re:Invent 2014 \(SEC404\): Incident Response in the Cloud](#)
- [AWS re:Invent 2015 \(SEC308\): Wrangling Security Events in The Cloud](#)
- [AWS re:Invent 2015 \(SEC316\): Harden Your Architecture with Security Incident Response Simulations](#)
- [AWS re:Invent 2016 \(SEC313\): Automating Security Event Response, from Idea to Code to Execution](#)
- [AWS re:Invent 2017 \(SID302\): Force Multiply Your Security Team with Automation and Alexa](#)

- [AWS re:Invent 2016 \(SAC316\): Security Automation: Spend Less Time Securing Your Applications](#)
- [AWS re:Invent 2016 \(SAC304\): Predictive Security: Using Big Data to Fortify Your Defenses](#)
- [AWS re:Invent 2017 \(SID325\): Amazon Macie: Data Visibility Powered by Machine Learning for Security and Compliance Workloads](#)
- [AWS London Summit 2018: Automating Incident Response and Forensics in AWS](#)

Third-Party Tools

The following links to third-party tools are external and are not endorsed by AWS. AWS offers no guarantees or representations of any kind about these tools or pages.

- [AWS IR](#) – Python installable command line utility for mitigation of host and key compromises
- [MargaritaShotgun](#) – Remote Memory Acquisition Tool
- [ThreatPrep](#) – Python module for evaluation of AWS account best practices around incident handling readiness
- [ThreatResponse Web](#) – Web based analysis platform for use with the AWS_IR command line tool.
- [GRR Rapid Response](#) – Remote live forensics for incident response
- [Linux Write Blocker](#) – The kernel patch and user-space tools to enable Linux software write blocking

Industry References

- [NIST SP 800-61R2: Computer Security Incident Handling Guide](#)

Document Revisions

Date	Description
June 2019	First publication

Appendix A – Examples

Example AWS CloudTrail Event

The following example shows that an IAM user named Alice used the AWS CLI to call the Amazon EC2 `StopInstances` action by using `ec2-stop-instances`.

```
{
  "Records": [
    {
      "eventVersion": "1.0",
      "userIdentity": {
        "type": "IAMUser",
        "principalId": "EX_PRINCIPAL_ID",
        "arn": "arn:aws:iam::123456789012:user/Alice",
        "accountId": "123456789012",
        "accessKeyId": "EXAMPLE_KEY_ID",
        "userName": "Alice"
      },
      "eventTime": "2014-03-06T21:01:59Z",
      "eventSource": "ec2.amazonaws.com",
      "eventName": "StopInstances",
      "awsRegion": "us-east-2",
      "sourceIPAddress": "205.251.233.176",
      "userAgent": "ec2-api-tools 1.6.12.2",
      "requestParameters": {
        "instancesSet": {
          "items": [
            {
              "instanceId": "i-ebeaf9e2"
            }
          ]
        },
        "force": false
      },
      "responseElements": {
        "instancesSet": {
          "items": [
            {
              "instanceId": "i-ebeaf9e2",
              "currentState": {
                "code": 64,
                "name": "stopping"
              },
              "previousState": {
                "code": 16,
                "name": "running"
              }
            }
          ]
        }
      }
    }
  ]
}
```

Example AWS CloudWatch Event

The following Amazon CloudWatch Event example shows that an AWS IAM user named `jane-roe-test` was found publicly exposed on `www.github.com`, and could be abused by unauthorized users.

```
{
  "check-name": "Exposed Access Keys",
  "check-item-detail": {
    "Case ID": "02648f3b-e18f-4019-8d68-ce25efe080ff",
    "Usage (USD per Day)": "0",
    "User Name (IAM or Root)": "jane-roe-test",
    "Deadline": "1440453299248",
    "Access Key ID": "AKIAIOSFODNN7EXAMPLE",
    "Time Updated": "1440021299248",
    "Fraud Type": "Exposed",
    "Location": "www.github.com"
  },
  "status": "ERROR",
  "resource_id": "",
  "uuid": "cce6d28f-e44b-4e61-aba1-5b4af96a0f59"
}
```

Example Infrastructure Domain CLI Activities

The following AWS CLI commands show an example of responding to an event within the infrastructure domain. This example uses the AWS APIs to perform many of the initial incident response activities described in this paper.

```
# Anomaly detected on IP X.X.X.X. Capture that instance's metadata
> aws ec2 describe-instances --filters "Name=ip-
address,Values=X.X.X.X"
```

```
# Protect that instance from accidental termination
> aws ec2 modify-instance-attribute --instance-id i-abcd1234 --
attribute disableApiTermination --value true
```



```
# Switch the EC2 instance's Security Group to a restricted Security Group
> aws ec2 modify-instance-attribute --instance-id i-abcd1234 --groups sg-a1b2c3d4
```

```
# Detach from the Auto Scaling Group
> aws autoscaling detach-instances --instance-ids i-abcd1234 --auto-scaling-group-name web-asg
```

```
# Deregister the instance from the Elastic Load Balancer
> aws elb deregister-instances-from-load-balancer --instances i-abcd1234 --load-balancer-name web-load-balancer
```

```
# Create an EBS snapshot
> aws ec2 create-snapshot --volume vol-12xxxx78 --description "ResponderName-Date-REFERENCE-ID"
```

```
# Create a new EC2 instance from the Forensic Workstation AMI
> aws ec2 run-instances --image-id ami-4n6x4n6x --count 1 --instance-type c4.8xlarge --key-name forensicPublicKey --security-group-ids sg-1a2b3c4d --subnet-id subnet-6e7f819e
```

```
# Create a new EBS volume copy from the EBS snapshot
> aws ec2 create-volume --region us-east-1 --availability-zone us-east-1a --snapshot-id snap-abcd1234 --volume-type io1 --iops 10000
```

```
# Attach the volume to the forensic workstation
> aws ec2 attach-volume --volume-id vol-1234abcd --instance-id i-
new4n6x --device /dev/sdf
```

```
# Create a security group rule to allow the new Forensic
Workstation to communicate to the contaminated instance.
> aws ec2 authorize-security-group-ingress --group-id sg-a1b2c3d4 -
-protocol tcp --port 0-65535 --source-group sg-1a2b3c4d
```

```
# Tag the contaminated instance with the ticket or reference ID
> aws ec2 create-tags -resources i-abcd1234 -tags
Key=Environment,Value=Quarantine:REFERENCE-ID
```

Appendix B – Cloud Capability Definitions

Amazon Web Services offers over 150 cloud services and thousands of features. Many of these provide native detective, preventative, and responsive capabilities, and others can be used to architect custom security solutions. This section includes a subset of those services that are most relevant to incident response in the cloud.

Logging and Events

AWS CloudTrail – AWS CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of your AWS account. With CloudTrail, you can log, continuously monitor, and retain account activity related to actions across your AWS infrastructure. 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. This event history simplifies security analysis, resource change tracking, and troubleshooting.

Validated log files are invaluable in security and forensic investigations. To determine whether a log file was modified, deleted, or unchanged after CloudTrail delivered it, you can use CloudTrail log file integrity validation. This feature is built-in using industry standard algorithms: SHA-256 for hashing and SHA-256 with RSA for digital signing. This makes it computationally infeasible to modify, delete or forge CloudTrail log files without detection.

By default, the log files delivered by CloudTrail to your bucket are encrypted by Amazon server-side encryption. You can optionally use the AWS Key Management Service (KMS) managed keys (SSE-KMS) for your CloudTrail log files.

Amazon CloudWatch Events – Amazon CloudWatch Events delivers a near real-time stream of system events that describe changes in AWS resources, or when API calls are published by AWS CloudTrail. Using simple rules that you can quickly set up, you can match events and route them to one or more target functions or streams. CloudWatch Events becomes aware of operational changes as they occur. CloudWatch Events can respond to these operational changes and take corrective action as necessary, by sending messages to respond to the environment, activating functions, making changes, and capturing state information. Some security services, such as Amazon GuardDuty, produce their output in the form of CloudWatch Events.

AWS Config – AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. Config continuously monitors and records

your AWS resource configurations and enables you to automate the evaluation of recorded configurations against desired configurations. With Config, you can review changes in configurations and relationships between AWS resources, manually or automatically. You can review detailed resource configuration histories, and determine your overall compliance against the configurations specified in your internal guidelines. This enables you to simplify compliance auditing, security analysis, change management, and operational troubleshooting.

Amazon S3 Access Logs – If you store sensitive information in an Amazon S3 bucket, you can enable S3 access logs to record every upload, download, and modification to that data. This log is separate from, and in addition to, the CloudTrail logs that record changes to the bucket itself (such as changing access policies and lifecycle policies).

Amazon CloudWatch Logs – You can use Amazon CloudWatch Logs to monitor, store, and access your log files (such as your operating system, application, and custom log files) from your Amazon Elastic Compute Cloud (Amazon EC2) instances using the CloudWatch Logs agent. Additionally, Amazon CloudWatch Logs can capture logs from AWS CloudTrail, Amazon Route 53 DNS Queries, VPC Flow Logs, Lambda functions, and other sources. You can then retrieve the associated log data from CloudWatch Logs.

Amazon VPC Flow Logs – VPC flow logs enable you to capture information about the IP traffic going to and from network interfaces in your VPC. After you've created a flow log, you can view and retrieve its data in Amazon CloudWatch Logs. VPC flow logs can help you with a number of tasks. For example, you can use flow logs to troubleshoot why specific traffic is not reaching an instance, which can help you diagnose overly restrictive security group rules. You can also use flow logs as a security tool to monitor the traffic to your instance.

AWS WAF Logs – AWS WAF now supports full logging of all web requests that are inspected by the service. You can store these logs in Amazon S3 for compliance and auditing needs, as well as to use them for debugging and additional forensics. These logs help you to understand why certain rules are triggered and why certain web requests are blocked. You can also integrate the logs with your SIEM and log analysis tools.

Other AWS Logs – With the pace of innovation, we continue to deploy new features and capabilities for customers practically every day, which means that there are dozens of AWS services that provide logging and monitoring capabilities. For information about the features available for each AWS service, see the AWS documentation for that service.

Visibility and Alerting

AWS Security Hub – AWS Security Hub gives you a comprehensive view of your high-priority security alerts and compliance status across AWS accounts. With Security Hub, you have a single place that aggregates, organizes, and prioritizes your security alerts or findings from multiple AWS services, such as Amazon GuardDuty, Amazon Inspector, and Amazon Macie, as well as from AWS Partner solutions. Your findings are visually summarized on integrated dashboards with actionable graphs and tables. You can also continuously monitor your environment using automated compliance checks based on the AWS best practices and industry standards your organization follows.

Amazon GuardDuty – Amazon GuardDuty is a managed threat detection service that continuously monitors for malicious or unauthorized behavior to help you protect your AWS accounts and workloads. It monitors for activity such as unusual API calls or potentially unauthorized deployments that indicate a possible account compromise. GuardDuty also detects potentially compromised instances or reconnaissance by attackers.

GuardDuty identifies suspected attackers through integrated threat intelligence feeds and uses machine learning to detect anomalies in account and workload activity. When a potential threat is detected, the service delivers a detailed security alert to the GuardDuty console and AWS CloudWatch Events. This makes alerts actionable and easy to integrate into existing event management and workflow systems.

Amazon Macie – Amazon Macie is an AI-powered security service that helps you prevent data loss by automatically discovering, classifying, and protecting sensitive data stored in AWS. Amazon Macie uses machine learning to recognize sensitive data such as personally identifiable information (PII) or intellectual property, assigns a business value, and provides visibility into where this data is stored and how it is being used in your organization. Amazon Macie continuously monitors data access activity for anomalies, and delivers alerts when it detects a risk of unauthorized access or inadvertent data leaks.

AWS Config Rules – An AWS Config Rule represents the preferred configurations for a resource and is evaluated against configuration changes on the relevant resources, as recorded by AWS Config. You can see the results of evaluating a rule against the configuration of a resource on a dashboard. Using Config Rules, you can assess your overall compliance and risk status from a configuration perspective, view compliance trends over time, and find which configuration change caused a resource to be out of compliance with a rule.

AWS Trusted Advisor – AWS Trusted Advisor is an online resource to help you reduce cost, increase performance, and improve security by optimizing your AWS environment. Trusted Advisor provides real time guidance to help you provision your resources following AWS best practices. The full set of Trusted Advisor checks, including CloudWatch Events integration, is available to Business and Enterprise support plan customers.

Amazon CloudWatch – Amazon CloudWatch is a monitoring service for AWS Cloud resources and the applications you run on AWS. You can use Amazon CloudWatch to collect and track metrics, collect and monitor log files, set alarms, and automatically react to changes in your AWS resources. Amazon CloudWatch can monitor AWS resources, such as Amazon EC2 instances, Amazon DynamoDB tables, and Amazon RDS DB instances, as well as custom metrics generated by your applications and services, and any log files your applications generate. You can use Amazon CloudWatch to get system-wide visibility into resource utilization, application performance, and operational health. You can use these insights to react and keep your application running smoothly.

AWS Inspector – Amazon Inspector is an automated security assessment service that helps improve the security and compliance of applications deployed on AWS. Amazon Inspector automatically assesses applications for vulnerabilities or deviations from best practices. After performing an assessment, Amazon Inspector produces a detailed list of security findings prioritized by level of severity. These findings can be reviewed directly or as part of detailed assessment reports which are available through the Amazon Inspector console or API.

Automation

AWS Lambda – AWS Lambda is a serverless compute service that runs your code in response to events and automatically manages the underlying compute resources for you. You can use AWS Lambda to extend other AWS services with custom logic, or create your own back-end services that operate at AWS scale, performance, and security. Lambda runs your code on high-availability compute infrastructure and performs all the administration of the compute resources for you. This includes server and operating system maintenance, capacity provisioning and automatic scaling, code and security patch deployment, and code monitoring and logging. All you have to do is supply the code.

AWS Step Functions – AWS Step Functions makes it easy to coordinate the components of distributed applications and microservices using visual workflows. Step

Functions provides a graphical console to arrange and visualize the components of your application as a series of steps. This makes it simple to build and run multistep applications. Step Functions automatically triggers and tracks each step, and retries when there are errors, so your application executes in order and as expected.

Step Functions logs the state of each step, so when things do go wrong, you can diagnose and debug problems quickly. You can change and add steps without writing code, so you can easily evolve your application and innovate faster. AWS Step Functions is part of the AWS Serverless Platform, and makes it simple to orchestrate AWS Lambda functions for serverless applications. You can also use Step Functions for microservices orchestration using compute resources such as Amazon EC2 and Amazon ECS.

AWS Systems Manager – AWS Systems Manager gives you visibility and control of your infrastructure on AWS. Systems Manager provides a unified user interface so you can view operational data from multiple AWS services and enables you to automate operational tasks across your AWS resources. With Systems Manager, you can group resources by application, view operational data for monitoring and troubleshooting, and take action on your groups of resources. Systems Manager can keep your instances in their defined state, perform on-demand changes, such as updating applications or running shell scripts, and perform other automation and patching tasks.

Secure Storage

Amazon S3 – Amazon S3 is object storage built to store and retrieve any amount of data from anywhere. It is designed to deliver 99.999999999% durability, and stores data for millions of applications used by market leaders in every industry. Amazon S3 provides comprehensive security and compliance capabilities that meet the most stringent regulatory requirements. It gives customers flexibility in the methods they use to manage data for cost optimization, access control, and compliance. Amazon S3 provides query-in-place functionality, which enables you to run powerful analytics directly on your data at rest in Amazon S3. Amazon S3 is the most supported cloud storage service available, with integration from the largest community of third-party solutions, systems integrator partners, and other AWS services.

Amazon S3 Glacier – Amazon S3 Glacier is a secure, durable, and extremely low-cost cloud storage service for data archiving and long-term backup. It is designed to deliver 99.999999999% durability, and provides comprehensive security and compliance capabilities that can help meet the most stringent regulatory requirements. Amazon S3 Glacier provides query-in-place functionality, which enables you to run powerful

analytics directly on your archive data at rest. To keep costs low yet suitable for varying retrieval needs, Amazon S3 Glacier provides three options for access to archives, from a few minutes to several hours.

Custom

The aforementioned services and features are not an exhaustive list. Amazon Web Services is continuously adding new capabilities. For more information, we encourage you to review the [AWS What's New](#) and [AWS Security](#) pages. In addition to the security services that AWS offers as native cloud services, you might be interested in building your own capabilities on top of AWS services.

Although we recommend enabling a base set of security services within your accounts, such as AWS CloudTrail, Amazon GuardDuty, and Amazon Macie, you might eventually want to extend these capabilities to derive additional value from your log assets. There are a number of partner tools available, such as those listed in our APN Security Competency program. You might also want to write your own queries to search your logs. With the extensive number of managed services that AWS offers, this has never been easier. There are many additional AWS services that can assist you with investigation that are outside the scope of this paper, such as Amazon Athena, Amazon Elasticsearch Service, Amazon QuickSight, Amazon Machine Learning, and Amazon EMR.