

FRANKKIMFRANKKIMFRANKKIMFRANKKIM
DAVESHACKLEFORDDAVESHACKLEFORD
JOHNPESCATOREJOHNPESCATOREJOHN
MATTBROMILEYMATTBROMILEYMATTB
KYLEEDICKINSONKYLEEDICKINSONKYLE
DAVIDHAZARDAVIDHAZARDAVIDHAZAR
JUSTINHENDERSONJUSTINHENDERSONJ
JOSHTHURSTONJOSHTHURSTONJOSH

SANS

Practical Guide to Security in the AWS Cloud

NAMLENAMLENAMLENAMLENAMLENAML
DAVIDSZILIDAVIDSZILIDAVIDSZILIDAV
SHAUNMCCULLOUGHSHAUNMCCULLOUG
NATHANGETTYNATHANGETTYNATHANG
KEVINGARVEYKEVINGARVEYKEVINGARV
BRIANRUSSELLBRIANRUSSELLBRIANRU
THOMASJ.BANASIKTHOMASJ.BANASIKTH
J.MICHAELBUTLERJ.MICHAELBUTLERJ.M
SOUNILYUSOUNILYUSOUNILYUSOUNILY

Table of Contents

Introduction

1. Cloud Security Overview – Frank Kim.....	5
2. What to Expect and Why it Matters – Rob Lee	8
3. A Practitioner’s Process to Mapping Frameworks and Standards to Technology – Josh Thurston	10
4. How to Optimize Security Operations in the Cloud Through the Lens of the NIST Framework – John Pescatore	18

Automating Compliance and Securing Data and Applications in AWS

5. How to Automate Compliance and Risk Management for Cloud Workloads – Matt Bromiley	38
6. How to Build a Data Security Strategy in AWS – Dave Shackleford	49
7. How to Design a Least Privilege Architecture in AWS – Dave Shackleford	65
8. How to Secure App Pipelines in AWS – Dave Shackleford	78
9. How to Protect a Modern Web Application in AWS – Shaun McCullough	96

Enhancing Protection of Applications, Devices, and Networks

10. How to Protect Enterprise Systems with Cloud-Based Firewalls – Kevin Garvey	114
11. How to Implement a Software-Defined Network Security Fabric in AWS – Dave Shackleford	127
12. How to Build an Endpoint Security Strategy in AWS – Thomas J. Banasik	142
13. How to Leverage a CASB for Your AWS Environment – Kyle Dickinson	151

Improving Visibility, Threat Detection, and Investigations in AWS

14. How to Build a Security Visibility Strategy in the Cloud – Dave Shackleford	160
15. How to Improve Security Visibility and Detection/Response Operations in AWS	173
– Dave Shackleford	
16. How to Build a Threat Detection Strategy in Amazon Web Services (AWS) – David Szili	189
17. How to Perform a Security Investigation in AWS – Kyle Dickinson	206
18. How to Leverage Endpoint Detection and Response (EDR) in AWS Investigations	216
– Justin Henderson	
19. How to Build a Threat Hunting Capability in AWS – Shaun McCullough	230

Solution Guidance in AWS

20. Solution Guidance for Application Security in AWS – Nathan Getty	252
21. Solution Guidance for Cloud-Based Firewalls in AWS – Brian Russell	265
22. Solution Guidance for Endpoint Security in AWS – David Hazar	280
23. Solution Guidance to Cloud Security Posture Management in AWS – Kyle Dickinson	303
24. Solution Guidance for SIEM in AWS – J. Michael Butler	312
25. How to develop a scalable security strategy in a multi-account environment	327
in AWS – Nam Le	

Prioritizing Security Controls in AWS

26. How to Prioritize Security Controls for Better Visibility and Context in AWS	334
– Sounil Yu	
27. How to Prioritize Security Controls for Sensitive Data and Applications in AWS	349
– Sounil Yu	

Conclusion: Sounil Yu	363
------------------------------------	-----

Chapter 1: Cloud Security Overview



Frank Kim

SANS Faculty Fellow and Curriculum Lead

"Cloud computing has become a major defining factor in the current and future state of information security, with the business reasons for moving to the cloud simply too overwhelming to ignore.

However, the cloud represents big change for almost all organizations, and security must be part of that evolution in order to succeed. In terms of industry momentum, we've now reached the point where every cybersecurity professional needs to be knowledgeable about the cloud to varying degrees.

As a security professional, you need to do three things in parallel:

- Understand how the major cloud providers work and the plenitude of services that they offer.*
- Understand the technical details of each platform to ensure that you have secured your specific implementation appropriately.*
- Ensure your teams transform the way they do their work in order to leverage cloud services and automation in a way that improves the effectiveness of security itself."*

This book provides you with a comprehensive collection of technical resources that you can use to arm yourself with the foundational knowledge required in today's cloud-first world.

Taken together, these resources model the whole life cycle of security, touching on aspects of the functions of the NIST Cybersecurity Framework—Identify, Protect, Detect and Respond.

This collection is a good place to start if you're looking to build out your cloud security knowledge base, because the technical detail provided in these reports and guides will enable you to start crafting a technical roadmap for your organization's transition to the cloud.

The reason I say that this is a good place to start, however, is that it's what you do next with the information you learn that matters most. Building and leading a cloud security program is not just about the technical controls; it's about the management, governance, people and process items as well. It's not just about implementing the right technology; it's also about the overall mission and vision of the organization.

So the question becomes, how do you align with that mission to ensure that you're achieving the larger business objectives in addition to your technical activities?

It might not be obvious, but the topics described in these resources are the foundational elements of your overall cloud security journey. Think of each resource as a piece of the puzzle that, once put together, creates a bigger picture. Now, it's up to you to connect the dots. As you read, I encourage you to challenge yourself to think about how these papers come together to create a broader view of the cloud. Doing so will enable you to build an overall cloud security roadmap for your business—not just a technical roadmap, but a business roadmap for the cloud.

It's a valuable exercise, to be sure, and it will make all the difference if you go into it with a strong understanding of your business objectives and drivers. With your business reasons for moving to the cloud top of mind, you'll be better able to lay out your objectives and roadmap to ensure that you accomplish what you need to in your first year and beyond.

It can be challenging to see how the day-to-day security activities discussed in these resources contribute to achieving your overall business goals, but you can treat this book as a checklist of sorts, and check things off in your mind as you read about the capabilities you need to implement in your organization. By doing so you will steadily improve the maturity of your overall cloud security program.

Just as the web has defined the previous 20 years of technology change, I believe that the cloud will be the defining element of the next 20 years. If you haven't already started building your cloud security knowledge and roadmap, there's no better time to start than now.

About the Author

Frank Kim leads the management and cloud security curricula for SANS, developing courses on strategic planning, leadership, DevSecOps and cloud security. He is also a SANS faculty fellow and author of MGT512, MGT514, and SEC540. Previously, Frank served as CISO at the SANS Institute, leading its information risk function, and was executive director of cybersecurity at Kaiser Permanente, where he built an innovative security program to serve one of the nation's largest not-for-profit health plans and integrated healthcare provider. Currently, as founder of ThinkSec, a security consulting and CISO advisory firm, Frank helps leaders develop business-driven security programs.

Chapter 2: What to Expect and Why it Matters



Rob Lee

SANS Fellow and Chief Curriculum Director and Faculty Lead

Section 1: Automating Compliance and Securing Data and Applications in AWS

"The rapid adoption of moving to cloud infrastructure across an organization has only accelerated recently, due to supporting remote workforce and customer needs universally. Addressing compliance in the face of such rapid change should ensure that critical data and applications remain secure across this technological shift."

Section 2: Enhancing Protection of Applications, Devices, and Networks

"Security should never be an afterthought, even in the cloud. Securing cloud data and capabilities is a critical step in proper deployments and should not be quickly implemented without adequate understanding. Cloud security mechanisms are not insurmountable and can be applied to ensure lower risk profiles and key monitoring to detect threats. This is the cloud equivalent of not skipping "leg day" at the gym—don't skip cloud security."

Section 3: Improving Visibility, Threat Detection, and Investigations in AWS

"Cloud threat hunting, detection, and mitigation should not be frustrating. Cloud capabilities to enable threat detection and response over the past few years have enabled organizations to get ahead of threats before they become the next headline. The challenge, though, has been educating security teams to understand how to leverage these new capabilities into their security operations capabilities. This section covers many areas that are now required reading for organizations to operationalize a proper detection, response, and mitigation capabilities across the cloud."

Section 4: Solution Guidance in AWS

"I always joke with my friends and family that I can fix anything with duct tape and YouTube. Understanding concepts and capabilities in cloud security are good to know. Having a useful how-to guide is key to solving quick problems and gaining experience effectively—just like having duct tape and YouTube. Having key cloud walk-throughs on where to begin is the best thing about this book."

About the Author

Rob graduated from the U.S. Air Force Academy and served in the U.S. Air Force as a founding member of the 609th Information Warfare Squadron, the first U.S. military operational unit focused on information operations. Later, he was a member of the Air Force Office of Special Investigations (AFOSI) where he led a team conducting computer crime investigations, incident response, and computer forensics. Prior to starting his own firm, he directly worked with a variety of government agencies, U.S. Department of Defense, and intelligence communities as the technical lead for a vulnerability discovery and an exploit development team, lead for a cyber forensics branch, and lead for a digital forensic and security software development team. Rob was also a director for MANDIANT, a company focused on investigating advanced adversaries, such as the APT, for five years prior to starting his own business.

Rob co-authored the book *Know Your Enemy*, 2nd Edition. Rob earned his MBA from Georgetown University in Washington DC. Rob is also a co-author of the MANDIANT threat intelligence report *M-Trends: The Advanced Persistent Threat*.

Chapter 3: A Practitioner's Process to Mapping Frameworks and Standards to Technology



Josh Thurston

Sr. Category Lead, Security at AWS

"The digital world has gravitated towards building a cybersecurity practice based on frameworks and standards. The number of frameworks has grown in volume, spanning industries including healthcare, financial services, retail, government, and just about everything in between. Framework adoption growth is driven by mandates such as the Payment Card Industry mandate for any organization processing payment cards (credit and debit cards) to be PCI-DSS compliant. Other frameworks are used as a matter of choice. Take the NIST Cybersecurity Framework as an example. U.S. President Donald Trump signed the Cybersecurity of Federal Networks Executive Order, which requires federal agencies to follow the NIST CSF. The NIST CSF is a mandate for federal agencies, while private sector organizations may choose to follow the NIST CSF.

To consolidate the usage of the terms "frameworks" and "standards," this chapter will combine them into a single term: "blueprint."

The ideology of using a blueprint has changed from a nice-to-have to an expectation. This movement stems from regulations and, in some cases, executive orders, as previously described. While the mindset of a buyer has shifted to blueprints, a challenge still exists. Industry blueprints are complex, and it can be challenging for buyers to know which products or services are available as they relate to a 'control' within the buyer's blueprint. To overcome this, security vendors began mapping their products and services to various blueprints. This can be an exhaustive exercise for sellers who ultimately want to meet buyer's needs, but the mappings

often do not perfectly align with the blueprints. Buyers and sellers both require an in-depth understanding of the blueprints available for use, and how to correctly map services to the blueprint controls.

The community needs a methodology to help individuals, organizations, and software vendors overcome the mapping challenge so that they are equipped to accurately map a blueprint to a product or service while understanding the true intent the blueprints are aiming to provide."

Practitioner's Process

The process may vary slightly between each of the blueprints, but this process can be followed to achieve the mission by using a small set of key primitives.

Primitive One

Read the control and ask: Is the control looking for People, Process, or Technology? For activities or actions that are manual, the answer is People. This can and should include people interfacing with technology, such as manually recording information on an asset. Process includes activities such as documentation, communications, escalations, analysis, and logic. Technology may be the most confusing of the three because it has become an integral part of people and process. Technology provides value to people by helping them scale, gain useful insights, automate processes, and accelerate work output.

An example to describe the difference between People, Process, and Technology

A bank has a datacenter with restricted physical access. Susan needs to enter the datacenter to update the firmware on several servers. When Susan approaches the entrance to the datacenter, she is greeted by a security guard. Susan signs the datacenter sign-in sheet, and the guard checks her identification and looks her up in the system to verify she has the authorization to enter. The security guard assigns an escort who accompanies Susan and monitors what servers she works on. This is all about People because the security guard is manually checking her identification and looking her up in the system.

Process comes into scope because the security guard has a checklist of tasks to complete for any datacenter visitor. The security guard validates identification, records the visitor with date and time, and finally assigns an escort who logs what servers the visitor gains physical access to. This process is set by the organization for historical reference. Very little technology has come into play, but the introduction of technology can make this example more modern. Susan approaches the datacenter and scans her

employee badge at the entry door. The back-end technology validates she is authorized to enter while recording her identity and the date and time of her entry. Technology can also track what server racks she swipes her badge at to create a trail of physical access to the racks. When Susan logs into any of the servers another trail of events is recorded. Technology may not have eliminated the security guard, but it may be possible to eliminate the need for the escort.

The remainder of the primitives and examples will focus on technology.

Primitive Two

The definition or true meaning of a control within the blueprint needs to be understood. To accurately identify a technology, ask the question: Does the primary function of the technology meet the control? Consolidating and packaging multiple capabilities into a single solution has become increasingly desirable and serves as a driving force to evolve technology. Vendors work continuously to provide more value in a solution while meeting customer consolidation demands. In other words, customers want to do more with less, and vendors want to streamline development and delivery.

Unfortunately, there is not a universal security solution that can deliver total security coverage. Products should have a primary function that is the basis for why a customer buys and uses it. Additional features are beneficial and provide value, but there is risk in delivering too many features with less quality. Outliers exist, but it is recommended to think of products for the primary function first when mapping.

Let's look at a mapping exercise using the NIST CSF Sub-category PR.DS-1: **Data at Rest is Protected**.

A product such as Data Encryption for Files and Folders protects data at rest, thus preventing unauthorized access.

The primary function of the Data Encryption for Files and Folders software is to protect data at rest and it meets the control requirement as described in the NIST CSF. Notice that the product is Data Encryption for Files and Folders, not the term Encryption alone. The term Encryption has too many options. People shop for products with a use case or a problem to solve in mind. One person may shop for encryption to encrypt data (files) while another person may shop for encryption to encrypt network traffic. The need to use very specific product categories or types is essential. Look back at the title of PR.DS-1. Now notice the three words that are in bold and recognize that this requires technology (primitive one), and the true meaning of the control is Protect + Data + [at] Rest. Primitives one and two have been achieved.

If someone rushed through the mapping exercise, they could do this incorrectly in several ways resulting

in mapping a data discovery tool, an email encryption solution, or even a SIEM, as explained in the next example.

Also, of note is that Data at Rest could define a couple of different things in that they may have varying solutions. Data at Rest could be interpreted as data on a drive, stored in something like an Amazon S3 bucket or a database, and even in a cloud storage and sync solution such as Dropbox. The compensating controls for each of those may vary from disk encryption, removable media encryption, or database encryption, and much more. Mapping seldom yields one solution, but the primitives hold up for every iteration.

Improper mapping

A product such as a SIEM solution can be used to log, aggregate, and correlate events related to data at rest. The primary function of a SIEM does not protect data at rest and does not meet the control requirement. A SIEM receives information (events) from various data sources in the environment that provide context (Who, What, When, Where, Why, How) to reveal unauthorized access and usage related to data at rest. Mapping a SIEM solution to PR.DS-1 would be incorrect because it cannot protect the data. All of the logs and events collection are after the fact. Primitive one is achieved by identifying Technology, but the second primitive is failed because the true definition was not matched properly with the technology capability based on a single word.

Primitive Three

Is the mapping believable? The second primitive is often difficult and takes a lot of time to complete. This third primitive is equally important because products are typically purpose-built. The notion that a product can identify, protect, detect, respond, and recover is false for most technologies today.

Let us look at a new example using the Functions of the NIST CSF and endpoint protection technology. Essentially, we are working the opposite direction as the previous examples where we looked at the control and mapped technology to it. Now we will look at a technology and break it apart to map it to controls.

A modern-day endpoint protection technology typically cannot identify, protect, detect, respond, and recover against vulnerabilities, threats, and exploits. The very name of the technology group signals the primary function, Protection. To better understand this concept, it is important to dig deeper into each of the five functions. Then we can accurately map to controls by narrowing the scope of the functions to examine and use the primitives for mapping.

The five functions of the CSF

This first of five functions in the CSF, Identify, is speaking to the identification and inventory of assets within an organization. In addition to assets, Identify also includes the inventory and identification of threats in the world. Both Inventory and Identification are not primary functions for most endpoint protection products. Those products are typically using one or more backend services such as signatures, reputation lookup services, or machine learning (ML) and artificial intelligence (AI). Those backend services are designed to identify and inventory threats AND provide the inventory with threat information to endpoint technologies. The inventory and threat information doesn't have to be part of the endpoint protection product either. In many of the solutions available today, that information can be available as a quick reputation lookup. Endpoint protection products extract information such as MD5 hash and check the reputation of the file in a service. Lastly, ML and AI technology typically do not have threat information baked in. These offerings are most often created to quickly examine attributes and even behaviors to determine if the sample is good or bad rather than referencing a signature or even a reputation service. The cybersecurity ML/AI technology space is complex and requires a separate discussion. For simplicity purposes, note that they do not have an inventory capability.

The Protect function of the CSF is speaking to the primary function of many endpoint technologies. These technologies are deployed and configured to protect and reduce risk. This is achieved by reducing access and eliminating the ability to exploit a vulnerability. Evidence is seen in the form of blocking an unauthorized attempt to access a service, registry key, port, credentials, etc. The key to success or failure is the configuration and the upstream technology mapped to the Identify function. If an endpoint protection product is configured improperly, it will produce a false negative or false positive. The product may miss something it was intended to stop, and the product may stop something it was intended to allow. Configuration is crucial to success and depends on the organizational risk tolerance (too tight vs. too lenient). This balancing act is a driving force for years of innovation within the endpoint protection space resulting in signature-less, ML/AI products, sandboxing, and even reputation services to integrate with.

The Detect function of the CSF aligns to a subset of endpoint technologies, including but not limited to Endpoint Detection and Response (EDR). EDR solutions came to fruition because endpoint protection solutions are not perfect. Not only do endpoint protection products miss threats from time to time, they often collect a different type of data compared to EDR products. If an endpoint technology fails to protect an asset, it is typically because it was:

1. Configured incorrectly

2. Not capable of protecting a vulnerability (known or unknown)

3. Not informed of the threat from an upstream service (DAT, Reputation, ML/AI)

In these three scenarios where an endpoint technology creates an event for a successful exploit, it essentially means that Protect failed. Failing to protect while logging the event does not mean that the product is meeting the Detect function because the product is technically logging security events, not a threat.

Some individuals experience cognitive dissonance because they are trained to think that endpoint protection products detect threats. This is false when the product failed the primary function of detecting the threat based on unique identifiable attributes. Endpoint protection products leverage signatures and reputation lookup services while monitoring for suspicious behavior associated to a known vulnerability. If and when suspicious behaviors occur, the product alerts on the behavior but has no knowledge of the threat. In that moment, the separation of failures occurs. The product is revealing an incorrect A) configuration where the product was set to allow a threat, B) failure to protect a known vulnerability, or C) the failure to source information about a known threat.

Another issue with Detect is false positives. Often technology will match a behavior and create an alert in the absence of threat information.

A supporting example:

- An endpoint technology may be configured to protect a known threat that sends mass emails out from an endpoint. If the product is truly protecting, it will recognize the unique identifiable attributes and stop it.
- If the product alerts because a mass email was sent, the alert is based on the behavior. In this case, the product may false-positive when an actual user sends an email to a large number of recipients because it recognized the behavior in the absence of a threat.

The Respond function is once again a very small subset of endpoint technology. If a technology fails to identify a threat, fails to protect an asset vulnerability from a known threat, and properly detects an attempt on an unknown vulnerability then it will, in most scenarios, not be equipped to automatically respond. The Respond function requires human decision making and/or configuration.

The small subset of products that are truly Respond technologies are built with ML/AI to rapidly or

automatically make human decisions as a means to meet scale and scope. Using the EDR product category as an example should hold up. EDR products do not protect or reduce security risk, but they do have the capability to detect suspicious behaviors. But where does Response fit in? An EDR administrator or user can create rules in an EDR product for a response capability. Often an administrator recognizes a set of behaviors that are not desirable within the organization. The behaviors may be based on company policy or they are known malicious behaviors. Perhaps an organization does not want PowerShell to launch a specific behavior. These behaviors can be blocked (Response) automatically by configuration. Users can also learn of behaviors and leverage EDR to manually issue a response, such as terminate process or delete a file.

Another area where cognitive dissonance sets in relates to automation technologies such as Security Orchestration Automation Response (SOAR). These are labeled as Response technologies, but they have two fallacies.

1. SOAR products typically do not directly respond to the event. They send an action or command via integration to another product to execute the response. Ex. SOAR sends a command to an endpoint technology to terminate a process in an automated rule or configuration within the SOAR solution.
2. SOAR products require a human to create the workflow, aka Playbook, that includes the steps to send a respond command. These workflows include logic gates with Who, What, When, Where, Why, and How. Some SOAR products include thresholds and/or limits, but they seldom have prioritization and categorization for an asset, vulnerability, or threat.

Responding to a security event is process-driven because it requires a significant amount of log and analysis before the response activity is carried out. Respond technologies are highly desirable because they can reduce significant amounts of time and resources once they are configured to execute a set of actions based on human intellect and experience. More simply put, they are used to automate repetitive activities that humans have performed in the past.

The final function, Recover is perhaps the most dependent on People and Process. Very few Recover technologies are able to return an asset (device, application, network data, user) back to a known state of good health automatically or out-of-the-box. If the technology has this capability, it is most likely a micro Recover function because the technology does not facilitate learning and enhancing the rest of the environment to avoid security issues in the future.

This brings up the concept of “Pets vs. Cattle” (introduced by Bill Baker). At a very high level, Pets are personal assets that we care about deeply, assets like our laptops which are personal to us and hold sensitive data. We depend on these assets daily and fear their loss. Cattle are functional assets that serve a purpose, but we may have many of them for scalability and redundancy purposes. If something were to happen to a functional asset, we typically care less because we have multiple, or we can replace them with very little effort or impact.

Security practitioners unknowingly think about personal and functional assets as they approach recovery from a security event. When a functional asset is compromised, it may be very easy to re-deploy or even revert to a backup. Change procedures may be more relaxed for these assets because the risk and impact to the business are low. When a personal asset is compromised, it may be extremely difficult and require lengthy amounts of time to plan and execute a recovery. Security analysts will likely need to coordinate with asset owners and other teams within the organization to minimize data loss and disruption. The risk and impact may be very high. Recovery procedures for personal assets require heavy amounts of People and Process, but little to zero Technology.

Summary

Evaluate frameworks and standards using this information to build a defensible and operational security architecture. With this guidance and the use of the primitives, foundational and layered security technologies can be deployed in any environment to truly meet the control requirements. Know that vulnerability assessments and risk management is an iterative process that never ends. Technology should be accompanied by Process and People to be fully effective in utilizing a blueprint.

About the Author

Josh Thurston is a cyber security veteran and leader in driving company strategy and innovation. Over fifteen years, Thurston has helped organizations solve complex security challenges and mature their security programs. He has helped bring new innovative products to market, designed and managed Security Operations Centers, and advised organizations on architecture and strategy in the public and private sector. At AWS, he is responsible for several product categories in the AWS Marketplace Security catalog including SOAR, DFIR, Cloud Workload Security, and Cloud Posture Management.

Chapter 4: How to Optimize Security Operations in the Cloud Through the Lens of the NIST Framework



John Pescatore

SANS Director of Emerging Security Trends

"Before coming to SANS in 2013, I was the lead security analyst at Gartner for almost 14 years. In early 2000, I began to see enterprises moving rapidly to application service providers, which quickly turned into use of software-as-a service (SaaS) providers. By 2010, those early adopters were starting to use infrastructure-as-a-service (IaaS) offerings, first for test and dev environments and then for production workloads. Over that period I wrote a series of "Critical Security Questions to Ask..." research notes that served as checklists of the important security controls that should be in use by any well-run cloud-based service provider. They were extremely popular Gartner research notes, and as I talked with Gartner clients that were making secure and business-enabling use of the cloud, I identified a core set of processes in use by the leaders: attention to a foundation of basic security hygiene; a referenceable framework to base gap analyses against and to justify strategy and resource needs to management; and a focus on having an integrated approach to monitoring, protecting and restoring critical business capabilities and sensitive information across both on-premises and cloud-based resources.

Flash forward to 2019: The use of IaaS has become mainstream and key core security processes to focus on remain the same. This chapter focuses on using the NIST Cybersecurity Framework to make sure cloud-enabled business functions are at least as secure, and ultimately more secure, than they were before the availability of cloud services. Like everything we do at SANS, the goal is to provide security teams with actionable advice for supporting business goals with a secure approach to gaining the benefit of the cloud while avoiding or mitigating risk."

Introduction

The use of cloud services by businesses and government agencies has grown rapidly, with the movement of production workloads to infrastructure as a service (IaaS) growing at more than 35 percent per year.¹ This move to cloud-based services has required security programs to extend operations beyond the data center and to re-evaluate security architectures, processes and controls to maintain effectiveness and efficiency in their efforts to secure their sensitive business applications, be they local or cloud-based.

Some common success factors have emerged from enterprise cloud use cases where security has been maintained and even improved while moving critical services to IaaS:

- Integrate security services available from cloud service providers with third-party security products/services to secure business-critical cloud workloads. The virtualized infrastructure of IaaS offers native security services and capabilities that greatly reduce the attack aperture, and that can be augmented by additional third-party security controls when risk assessments require higher levels of protection.
- Extend security architecture, processes and controls across local data center applications and cloud IaaS implementations. Most enterprises use a mix of applications that run in local data centers, on external IaaS services and in hybrid configurations of both environments. Using common security controls and products across environments reduces the skills gap, eliminates data islands and silos, and makes it simpler to maintain a single security dashboard with a meaningful set of security metrics.
- Use an established framework to plan, implement and justify the changes needed to enable secure business use of IaaS. While securing cloud services relies on the same basic security ingredients used in traditional data center systems, the overall security architecture, processes and security controls must change to ensure that the necessary levels of reliability and safety are maintained. Basing the process on an established framework, such as the NIST Cyber Security Framework, ensures a thorough risk evaluation and implementation and provides a solid basis for justifying plans, strategies and resource requests to management.

Many businesses and government agencies have followed these guidelines to maintain their on-premises levels of security for production applications as those applications were moved to IaaS services. Even better, though, as new cloud security approaches emerged, they were able to raise the security level overall.

¹"IaaS Emerges as Fastest-growing Sector of the Global Public Cloud Market," ComputerWeekly, April 12, 2018, www.computerweekly.com/news/252438790/IaaS-emerges-as-fastest-growing-sector-of-the-global-public-cloud-market

“Today, organizations can build in security as an integrated part of the migration to IaaS services, optimizing security processes so they can be extended to work seamlessly across both local and external services.”

Keeping Business Safe—or Even Safer—in the Cloud

Cloud services security has evolved pretty much as security has evolved for all new technologies and innovations. Initially, security teams, with a healthy fear of the unknown, rated external cloud services as high risks because of reduced visibility and control, and so attempted to prevent their use. As the benefits of cloud services became apparent to business units and IT organizations, they adopted them, even if it meant bypassing the security organization. Security teams considered those cloud deployments to be rogue efforts, and therefore did not even evaluate the security arrangements.

In the face of security's resistance, CEOs began to tell CISOs, “We are moving to use cloud services, so tell us how to secure them or just get out of the way.” Only then did most security teams begin to try to reactively add security controls on top of cloud services and replicate on-premises data-centric security processes at virtualized cloud-based services. Their efforts did usually reduce risk, but at a high cost of business disruption. What's more, the tacked-on security processes were redundant and inefficient.

But things have improved. Today, organizations can build in security as an integrated part of the migration to IaaS services, optimizing security processes so they can be extended to work seamlessly across both local and external services. Similarly, security operations teams can focus on selecting products to implement security controls that are integrated across both environments, often minimizing vendor count, employee staffing and training requirements while enabling a single view of situational awareness and risk.

Differences in Securing Cloud Workloads

Just as any recipe for a meal can be broken down into the five basic tastes (sweet, sour, salty, bitter and umami), securing information always comes down to providing three basic security functions, the “CIA triad” of confidentiality, integrity and availability.² Security processes based on one or more of those basic functions deliver protect/detect/respond services using common security practices and products such as vulnerability assessment, configuration management, firewalls, anti-malware, SIEM and data protection.

All these security controls are necessary because of three key ongoing vulnerabilities:

- Applications and operating systems continue to have vulnerabilities that are not known until researchers find them and/or attackers exploit them.
- System administrators often make mistakes in configuring and maintaining servers and PCs.
- Users will always fall victim for scams such as phishing and malvertising.

The adoption of cloud services does not eliminate any of those areas of vulnerability—and can in fact magnify them, because the power of the cloud can greatly expand the vulnerabilities that result from weak practices in IT or security operations and administration.

On the other hand, IaaS brings the opportunity to significantly reduce the frequency of dangerous errors in operations and administration. The virtualized infrastructure of cloud services supports internal security mechanisms that evolving security processes can use in a number of ways:

- **Containers** — A container is a packaged unit of software that includes the application, the runtime operating systems, tools, libraries and so on.³ Well-prepared security teams can bake in configuration baselines and security agents that ensure that security controls will run anytime an application is launched.
- **Isolation** — Network segmentation has long been a proven way to limit exposure from attackers to an isolated segment and limit the spread of malware or other payloads. IaaS offerings can provide virtual private clouds that support segmentation at a granular level, with automated placement and enforcement when new servers are enabled. Containers also provide process isolation that enables CPU and memory utilization to be defined and limited on a granular basis.

²“Security Best Practices for IT Managers,” June 2013, www.sans.org/reading-room/whitepapers/bestprac/security-practices-project-managers-34257

³“Security Assurance of Docker Containers,” October 2016, www.sans.org/reading-room/whitepapers/assurance/security-assurance-docker-containers-37432

- **Orchestration and automation** — Many security processes are relatively static IF–THEN sequences that are often documented in playbooks. Orchestration defines the conditions and sequences, but implementation can be a highly manual process. Integration of security processes into cloud service management capabilities can automate many steps in security operations playbooks.

In this section we outlined the differences in securing cloud workloads. Next, we discuss using a security framework to address the needs security teams face.

The NIST Cyber Security Framework

The NIST Cyber Security Framework (CSF) came out of the Cybersecurity Enhancement Act of 2014,⁴ with the charter to be “a voluntary, consensus-based, industry-led set of standards, guidelines, best practices, methodologies, procedures, and processes to cost-effectively reduce cyber risks to critical infrastructure.”⁵ While there is nothing revolutionary about the NIST CSF, the “consensus-based, industry-led” approach resulted in widespread acceptance and adoption of the CSF by U.S. enterprises and the governments of several other countries.

The top level of the framework lists the five major functions (identify, protect, detect, respond and recover) of cybersecurity. These functions, which are intended to include all basic cybersecurity activities, are broken into 22 categories representing program-level outcomes required to maintain cybersecurity, as illustrated in Figure 1. These categories are further decomposed to list 98 subcategories that list specific results required to successfully implement the appropriate level of security.

“Securing information always comes down to providing three basic security functions, the “CIA triad” of confidentiality, integrity and availability.”


⁴National Institute of Standards and Technology, www.nist.gov/news-events/news/2018/04/nist-releases-version-1.1-its-popular-cybersecurity-framework

⁵Cybersecurity Enhancement Act of 2014, [www.congress.gov/bills/113th-congress/senate-bill/1353/text](http://www.congress.gov/bills/113/congress/senate/bills/1353/text)

The identify/protect/detect/respond/recover construct has proved to be a powerful tool in explaining to upper-level management the necessary core functions for protecting business systems, but in operational environments, very few processes or products perform just one of the top-level functions. For example, while firewalls are most closely identified with protective technology, they also play key roles in identify, protect, detect and respond. The construct also does not differentiate functional areas, processes and products that are important to use for proactive (before the attack) or reactive (during and after the attack) reduction of risk.

A more effective and efficient approach to selecting the most appropriate and effective security products and services to secure both data center and cloud-based systems is a scenario-based approach, which is covered in the next section.

Moving from Frameworks to Features, Talk to Walk



⁶"Introduction to the NIST CyberSecurity Framework for a Landscape of Cyber Menaces," Security Affairs, April 20, 2017, <https://securityaffairs.co/wordpress/58163/laws-and-regulations/nist-cybersecurity-framework-2.html>

Business units have been demanding the use of cloud-based services because of advantages they provide to efficiently deliver business services and adapt to changing needs. In order for security controls to be successful across both data center and cloud environments, security architectures, processes, controls and operations need to meet those same demands and provide the same seamless integration achievable in hybrid cloud services.

Why a Framework?

Regardless of the existing level of operations maturity, security teams face common needs:

- Adapting to changing business demands and evolving threats
- Obtaining management support for necessary resources and changes in IT or other areas
- Demonstrating improvement and providing risk assessment and forecasting
- Reducing the burden of satisfying auditors that security operations are compliant

A security framework, with its recommended set of security processes and controls, along with a risk assessment and management approach to match the appropriate set of controls to the business and threat environment, is an efficient way to meet these needs. Using an established framework can take the guesswork out of the process for smaller organizations, while allowing larger and more mature security operations to justify their decisions and resource requests to management and auditors.

Delivering Seamless Security Services

There are three key focus areas for delivering seamless security services across the data center and IaaS-based applications.

Integration of Infrastructure and External Security Controls at Each Boundary

Most organizations already have standard architectures for delivering identify/protect/detect/respond/restore services to data-center-based systems. When working with physical servers, organizations rely on a mix of security capabilities built into the Linux and Windows operating systems, as well as third-party host-based and network-based security controls. As local data centers moved to virtualization, another element was added to the mix: security primitives available in VMware or other underlying virtualization platforms. Similar, and often enhanced, security primitives are available from all major IaaS providers.

For companies other than startups, extending existing architectures to secure cloud-based services is the key first step. Those organizations should focus on integrating services at each boundary layer. See Figure 2.

In the early days of using the internet, many enterprises felt that there was a security gain by using products from different vendors at different layers in the architecture. However, real-world results proved this thinking to be false.⁷ For most security organizations, keeping the security architecture consistent across cloud services and the data center will support running the same security products across both environments. This will reduce training costs and administrative errors and also support more timely and accurate situational awareness and continuous monitoring.

Common Practice/Due Diligence Controls

Many security controls, such as firewalls, log monitoring and even intrusion detection systems, are mandated by compliance regimes (e.g., PCI DSS, HIPAA, FISMA, etc.) and represent due diligence controls. Any system containing sensitive or mission-critical data connected to the internet without a firewall and without log collection/monitoring/analysis would be considered noncompliant. While compliant does not always mean secure, noncompliant almost always represents unacceptable business risk.

⁷<https://www.gartner.com/document/500890?ref=solrResearch&refval=214539204&qid=d3f5b689a39463b6c77406155a9672a1>




Figure 2. Integrated Services at Each Boundary Layer

Best Practice/“Lean Forward Risk Reduction” Controls

As the continuing news of breaches makes clear, for many organizations “common practice” is insufficient to mitigate their actual. Best practice approaches that increase identify and protect levels and decrease time to detect, respond and restore are key, but require additional resources and skill levels. “Lean forward” organizations that have the staff skills and product/service budgets to deploy, tune and monitor advanced and proactive risk reduction controls generally are not the ones showing up in the breach headlines.

Using the NIST CSF Framework as a Starting Point for Putting Controls in Action

As mentioned earlier, the major security functions listed in the NIST CSF do not represent distinct product areas. However, Table 1 assigns a primary mapping for each major product area. This mapping can be used as a starting point in conjunction with a scenario-based approach to ensure that 1) you have no due diligence/compliance gaps, and 2) you have a solid baseline to which advanced capabilities can be added.

The decision on when to move beyond due diligence should be based on your own risk analysis. The NIST CSF points to the NIST Risk Management Framework,⁸ but many organizations have their own risk assessment and tracking processes that are outside the scope of this paper. The selection of architectures and products to implement security controls to protect cloud-based applications should be based on that assessment and the particular cloud deployment scenarios you face. The NIST CSF details the use of profiles and implementation tiers for this purpose. We will focus on a simplified approach based on the three most common cloud adoption scenarios facing businesses and government agencies:

- Dev/test environment
- Business app launched on or moved to IaaS
- Hybrid architecture

These scenarios represent the most frequent scenarios for securely moving business applications to cloud services in the typical order of adoption. While they do not represent every possible situation, these three scenarios generally provide a proven starting point you can tailor to your unique situation.

At the due diligence level, the basic security controls required are largely the same across the scenarios when business-critical or sensitive data is involved. The sections that follow describe the different drivers for each scenario with the assumption that such sensitive data is involved.

⁸Risk Management <https://www.gartner.com>

Table 1. Mapping Cloud Controls to the NIST CSF Framework

NIST CSF Functions		Primary Product Categories	
		Due Diligence	Advanced/Lean Forward
Proactive	Identify	Configuration management	AppSec testing
		System management	GRC
		Vulnerability assessment	Penetration testing
		Awareness training	
	Protect	Access management	Encryption
		Data masking	Intrusion prevention systems
		DDOS filtering	Secure image/container
		Endpoint protection	Strong authentication
		Firewall	Firewall policy management
		Ops skills training	
Reactive	Detect	Intrusion detection systems	Data analytics
		Network monitoring	Data loss prevention
		SIEM	
	Respond	Incident response services	Endpoint detect/respond
		Trouble ticket systems	Forensic analysis
	Recover	System/endpoint backup	High-avail/mirroring services

Dev/Test Environment

Moving a development and test environment to the cloud is often the first toe in the water for enterprise use of IaaS. The “pay as you go, not when you don’t need it” nature of IaaS is well-suited for this application. Rather than waste dedicated resources for development and test efforts that might only be used a small percentage of the time, an IaaS-based dev/test environment can be spun up and paid for only when actually needed.

All too often, the security organization is not involved in the migration, a circumstance with three downsides:

- Test data used in the IaaS instantiation often puts sensitive customer and business data at risk.
- That same environment can be used to rapidly evaluate operating systems and application patches, reducing exposure.
- The initial movement to dev/test on IaaS is an ideal chance for the security operation team to “plus up” its skills and develop knowledge around cloud capabilities and risks.

Data masking, obfuscation or encryption is a critical due-diligence requirement for dev/test environments. While realistic test data is necessary, you should never expose live customer data in dev/test usage. Similarly, standard boundary/perimeter network segmentation and monitoring as implemented by firewalls and IDS are required between this environment and the corporate network. If dev/test requires a live internet connection, the same controls are required at the internet connection side.

Because the entire purpose of a dev/test environment is to support an environment to deliver product-ready applications, the due diligence level includes application security (AppSec) testing tools/services that compliance regimes do not always require. Embedding AppSec testing into the development and test cycle is especially important in the rapid iteration cycles in agile/DevOps methodologies.

The traffic and user/endpoint behaviors on dev/test networks differ greatly from those on production systems, and advanced analytics and behavior-based detection/ prevention usually generate large volumes of false positives. With data masking in use, there is less of a need for data loss prevention, and dev/test environments generally do not require full DDoS protection. See Table 2.

Table 2. Security Control Set for Dev/Test Migration to IaaS

NIST CSF Functions		Primary Product Categories	
		Due Diligence	Advanced/Lean Forward
Proactive	Identify	AppSec testing	GRC
		Configuration management	Penetration testing
		System management	
		Vulnerability assessment	
	Protect	Access management	Encryption
		Data masking	Intrusion prevention systems
		Firewall	Secure image/container
		Ops skills training	Strong authentication
Reactive	Detect	Intrusion detection systems	
		SIEM	
	Respond	Incident response services	Endpoint detect/respond
		Trouble ticket systems	Forensic analysis
	Recover	System/endpoint backup	High-avail/mirroring services

Business App Launched on/Moved to IaaS

When a production application is launched from or moved to IaaS, the full range of confidentiality/integrity/availability services is required across all five NIST CSF functions to reach the due diligence level. From a product perspective, only data masking is typically not included in the architecture, because real product data is required and must be safeguarded. A typical example is a new web-based commerce application that will be first launched from an IaaS platform, but the same security principles apply to an existing application being updated and moved to IaaS.

The due diligence level of this scenario has two key goals:

- **Extend security configuration standards and continuous monitoring to IaaS.**

Every organization should have standards for the baseline configuration of all servers, applications, security controls and the like used in the production environment. These same standards, such as the Center for Internet Security Benchmarks,⁹ should be applied

⁹CIS Benchmarks, Center for Internet Security, www.cisecurity.org/cis-benchmarksdocument/500890

to applications running on IaaS. The processes for monitoring for misconfigurations and vulnerabilities should be identical for both data center applications and those running in IaaS. When it comes to product selection, it is key to have logging, monitoring and configuration/vulnerability analysis that integrates with a common SIEM platform and supports all applications.

- **Use common products for protect/detect infrastructure functions where possible.**

Most firewall, intrusion detection/protection, and endpoint protection products (and those like them) have both data center products and cloud-centric versions. Using the same vendor on IaaS as is used for data center security has all the advantages previously discussed.

When risk analysis requires higher levels of protection and resources (people, skills, budget) to support it, moving to the advanced security level generally means being proactive in avoiding or quickly mitigating vulnerabilities (AppSec testing, penetration testing); reducing unnecessary access privileges through secure access management, encryption and strong authentication (as a minimum for admin access); and reducing time to detect/respond/restore through the products and services listed.

In addition, you can raise the security bar for applications running on IaaS with such advanced cloud security capabilities as secure images and containers (discussed earlier). DDoS protection becomes more critical when an application is fully cloud-based. While cloud management platforms are not strictly security products, their use can increase the accuracy of asset management and vulnerability data, as well as support compliance reporting requirements. Governance, risk and compliance (GRC) platforms can greatly reduce the cost of demonstrating compliance (allowing more of the security budget to be focused on security), but they require large up-front investments in both procurement costs and administrative time and skills. See Table 3.

Hybrid Architecture

The final scenario is when organizations begin to run applications that span both local data centers and IaaS services in a near seamless manner. A common situation is expanding an application that has been running in a data center servicing one geographic region to global coverage using IaaS to expand capacity and proximity. The risk assessment used for the previous scenario ("Business App Launched on/Moved to IaaS") does not change for this scenario, but hybrid cloud environments do raise a number of unique challenges and opportunities:

- Changes in policy standards for identify and protect products must be distributed, validated and audited in an integrated manner across the environments.
- Detect products have a more complex environment to monitor, and behaviors in the more rigid data center environment often differ from what is seen on the IaaS environment.
- Forensic analysis as a respond function has more complicated attack paths to collect and analyze.
- If the IaaS environment supports a failover or mirroring capability, backup and recovery may be simplified in hybrid cloud environments.

Table 3. Security Control Set for Business App Launched on/Moved to IaaS

NIST CSF Functions		Primary Product Categories	
		Due Diligence	Advanced/Lean Forward
Proactive	Identify	Awareness training	AppSec testing
		Configuration management	GRC
		System management	Penetration testing
		Vulnerability assessment	Cloud management platforms
	Protect	Access management	Encryption
		DDOS filtering	Intrusion prevention systems
		Endpoint protection	Secure image/container
		Firewall	Strong authentication
		Ops skills training	Firewall policy management
Reactive	Detect	Intrusion detection systems	Data analytics
		Network monitoring	Data loss prevention
		SIEM	
	Respond	Incident response services	Endpoint detect/respond
		Trouble ticket systems	Forensic analysis
	Recover	System/endpoint backup	High-avail/mirroring services

For organizations that have not first moved through the first two scenarios, the migration to hybrid cloud services should not proceed without establishing a baseline of due diligence cloud infrastructure protection, monitoring and respond/restore capabilities, along with a security operations staff that has already expanded its skills to include cloud environments. From this starting point, staff can integrate the same advanced capabilities as in the previous scenario to raise security levels.

The primary difference in product selection for the hybrid cloud scenario is selecting products that you can deploy, manage and monitor across both environments (see Table 4). The typical starting point is to look at the security products in use on the data center side and see whether those vendors are listed in the IaaS provider's partners list or marketplace. Ideally you would use only products that are supported across the major IaaS providers, but there are simple workarounds for many product areas if you have to use different products:

- Network policy management tools support change control, auditing and analysis of firewall policies across multiple vendors.
- Any host-based product that supports syslog generation can report to a SIEM console.
- The output from disparate vulnerability assessment products that support the Security Content Automation Protocol (SCAP) can be consolidated by SIEM products.

Using Metrics to Assess and Communicate Effective Security Operations

From a security perspective, the movement to use IaaS does not change the need to collect meaningful security metrics. Metrics are needed not only to assess, evolve and optimize security operations, but also to provide accurate status, trend and risk data to management.

The minimal set of operations metrics that organizations should establish for their systems running on cloud services include:

- **Asset management accuracy** — What percentage of assets are identified and profiled correctly?
- **Time to detect** — How quickly is an attack detected?
- **Time to respond** — How quickly are incident response actions initiated?

- **Time to restore** — How quickly is incident response completed and full business services restored?
- **Real-time risk assessment** — What percentage of business-critical operations is currently at risk from known threats?

Table 4. Security Control Set for the Hybrid Cloud

NIST CSF Functions		Primary Product Categories	
		Due Diligence	Advanced/Lean Forward
Proactive	Identify	Configuration management	AppSec testing
		System management	GRC
		Vulnerability assessment	Penetration testing
		Awareness training	
	Protect	Access management	Encryption
		Data masking	Intrusion prevention systems
		DDOS filtering	Secure image/container
		Endpoint protection	Strong authentication
		Firewall	CASB
		Ops skills training	
Reactive	Detect	Intrusion detection systems	Data analytics
		Network monitoring	Data loss prevention
		SIEM	
	Respond	Incident response services	Endpoint detect/respond
		Trouble ticket systems	Forensic analysis
			Network policy management
	Recover	System/endpoint backup	High-avail/mirroring services

For most organizations, the metrics that security personnel show to CEOs and boards of directors will be different from operational metrics—the focus needs to be more strategic and show more connection to business services and less to attacks and threats. Figure 3 translates the key performance metrics into points that will resonate with CXOs and boards.

For most organizations, the metrics that security personnel show to CEOs and boards of directors will be different from operational metrics—the focus needs to be more strategic and show more connection to business services and less to attacks and threats. Figure 3 translates the key performance metrics into points that will resonate with CXOs and boards.

Translate Time to Detect and Time to Respond Improvement to:	
EFFICIENCY	EFFECTIVENESS
Decrease the cost of dealing with known threats.	Increase the speed of dealing with a new threat or technology.
Decrease the realized impact of residual risks.	Decrease the time required to secure a new business application, partner, or supplier.
Decrease the cost of demonstrating compliance.	Reduce incident cost: • Less downtime • Few customer defections
Increase incident count with constant staff.	Security as a competitive business factor
Maintain level of protection with less EBITDA impact.	

Figure 3. Connecting Metrics to Business Services

Summary

Thousands of businesses are successfully and safely using cloud services to meet business goals for increasing the agility and decreasing the cost of IT services. SANS has seen several common patterns across the security operations organizations that have been able to deliver the needed security architectures, processes and controls to enable safe business use of cloud services:

- Organizations use the NIST CSF Framework as a baseline and a tool to communicate and justify strategy, plans and resource needs to management.
- They involve the security team when IT first tries out IaaS, typically when dev/test is moved to the cloud. A robust selection of third-party security products in the cloud environment should be a key input into the evaluation of the IaaS provider.

- Teams extend the security architecture and processes to include applications running in the cloud, focusing on the most common business use cases.
- They maximize both effectiveness and efficiency by using the same third-party security products in the cloud that they use to secure on-premises applications (where possible).
- Once a secure baseline has been established for security operations in the cloud, security teams investigate cloud-specific security processes and controls that can result in advances over existing security practices.

Security teams will need to use mixes of people, processes and technologies to make sure business use of cloud services is secure. These patterns apply across all three of those areas. An honest assessment of your security operations team skills and processes completeness against the NIST CSF will enable you to evolve and extend security operations to enable business services while justifying needed changes and resources allocations.

About the Author

John Pescatore joined SANS as director of emerging technologies in January 2013 after more than 13 years as lead security analyst for Gartner, running consulting groups at Trusted Information Systems and Entrust, 11 years with GTE, and service with both the National Security Agency, where he designed secure voice systems, and the U.S. Secret Service, where he developed secure communications and surveillance systems “and the occasional ballistic armor installation.” John has testified before Congress about cybersecurity, was named one of the 15 most-influential people in security in 2008 and is an NSA-certified cryptologic engineer.

Chapter 5: How to Automate Compliance and Risk Management for Cloud Workloads



Matt Bromiley

SANS Certified Instructor

"As more and more services move to the cloud, organizations must be careful about the impact of this movement. Compliance is one area where many organizations may believe that the cloud offloads their responsibility-but nothing could be further from the truth. "Someone else's server" does not equate to "someone else's problem." Organizations must still be vigilant in protecting their data and their customers, and ensure that they remain compliant with all necessary regulations. In this chapter, I explore what compliance in the cloud means and the key things you need to keep when transitioning some of your services to a third party."

Introduction

There seems to be a constant battle between how fast businesses can grow and whether they can secure their customers' data. Many organizations get so wrapped up in trying to expand and scale for customer access that they make quick-fire, ad hoc decisions that negatively impact the security of the data of those very same customers. Complicating matters, the explosion of cloud-based services and offerings has led many organizations to quickly adopt services whose risks, quite frankly, they may not understand.

Of course, various compliance standards, such as PCI DSS and FedRAMP, have been developed to help organizations establish models to combat the loose handling of customer data. But this is where many organizations get lost. At the mere mention of the word "compliance," business and process owners tend to sink into an endless stream of acronym soup and never come up for air. But they do not have to fight this battle alone. The cloud is easy to deploy, but so is compliance.

While moving various elements of your business to the cloud does not remove the need for compliance, it does shape how you view, apply and assess compliance and risk management. In this paper, we focus on how moving to the cloud presents new compliance opportunities and how to seize them for your organization. We also examine a case study where a business has made a sudden shift to the cloud and look at some of the additional risk considerations it needs to make.

Last, we ask you to consider what may be a potential paradigm shift in how your organization approaches compliance and data security. In what we are calling "compliance-forward cloud planning," we encourage organizations to rethink the way they plan and deploy their cloud infrastructures, with compliance a focus from the beginning and not an afterthought. By focusing on compliance at the onset, organizations can make infrastructure decisions that will maintain compliance—not violate it.

Of course, if your organization has already moved to the cloud, compliance-forward planning may not be applicable, but the concepts pertaining to how to remain compliant certainly are. At the end of this paper, we hope you have some new thoughts and insight to bring to your team to discuss compliance and risk management options.

"Compliance-forward cloud planning is the concept of making cloud infrastructure planning decisions based on adhering to compliance of data first—not as an afterthought."

Risk Management: Protecting Your Customers

Before we discuss techniques to secure your data and infrastructure within the cloud, it is important to understand how your risk model changes within the cloud. Some organizations think that because the data exists in the cloud—that is, on someone else's system—compliance is no longer their responsibility. This is not the case, and such an assumption is likely to put a business at risk.

When an organization takes advantage of services and/or infrastructure within the cloud, only a handful of responsibilities transfer to the cloud provider. For example, the cloud provider is responsible for ensuring that the network and hardware remain up and functional. However, the organization is still responsible for the security of the data that is placed within the cloud resources. Figure 1 illustrates the division of responsibilities between an organization and its cloud provider.




Figure 1. Respective Responsibilities of Cloud Provider and Customer¹

¹ Note that your cloud provider may offer its own shared responsibility model. Check with your provider to verify what it does and does not provide.

Breaking Out of “Compliance-Backward”

As mentioned earlier, one of our goals with this paper is to shift to a compliance-forward frame of mind, where compliance becomes part of the design, not an afterthought or a nuisance. However, moving away from a compliance-backward approach is much easier said than done. Understanding compliance and what your data may be subject to can sometimes seem like a daunting task. In this section, we discuss common compliance standards and how to bring them into your organization.

Common Compliance Standards

Table 1 lists some of the more common compliance standards that your organization may encounter.

Bringing Compliance into Your Infrastructure

Whichever compliance standards your data may be subject to, cloud infrastructure provides multiple ways to achieve compliance. One of the most apparent is the ability to make use of multiple third-party vendors. Furthermore, your cloud provider may offer native, compliant-ready solutions that, when coupled with third-party integrations, can alleviate a lot of compliance headaches. Oftentimes, cloud providers facilitate third-party integrations and automations that allow for various application and infrastructure testing. Compliance is no different. Figure 2 describes techniques that you can use today to ensure your business remains compliant.

Table 1. High-Level Compliance Standards	
Standard	What It Protects or Defines
FedRAMP	The approach for security assessment and monitoring that must be in place to provide services to the U.S. government
HIPAA/ HITECH	Standards for securing the privacy of protected health information (PHI)
ISO 27001	Standards for security management and program implementation
PCI DSS	Payment cardholder data (CHD) or data used in transaction authorization (sensitive authorization data)




Figure 2. Techniques for Achieving Compliance

Case Study: Protecting Data on Multiple Angles

In recent years placing customer data within NoSQL and key-value databases has been a common strategy. With an easy-to-manage back end and rich front-end development options, NoSQL databases provide developers an easy-to-consume data format to enhance the customer experience. However, faster, compliance-second development also allows for compliance mishaps.

Let's examine an organization called "Bobby's Bits," a company that recently moved its infrastructure to the cloud. The following sections describe specific areas where compliance mishaps may negatively impact the organization and how to potentially mitigate or implement better controls.

Bobby's Bits: A New Cloud-Based Model

Bobby's Bits, a fictional organization, helps small businesses accept and process payments for online and in-store orders using payment methods other than cash or credit card. Bobby's business used to be fairly local, but because of some recent word-of-mouth marketing, the business has grown quite significantly. As a result, the owner had to hire a handful of developers and move his business away from the servers in his garage to the cloud. This move provided Bobby and the team easier access to all of the business's resources and allows them to automatically scale for busy days. Figure 3 shows a high-level diagram of the new business structure.

With this new cloud-based model, the company can scale quickly to meet the demands of its customers—and the demands of its customers' customers. But this rate of growth could be creating a compliance risk. In the next section, we examine a few areas where Bobby should probably exercise caution and slow down (and where you should too!).




Figure 3. High-Level Diagram of the New Cloud-Based Model

Protecting PCI Data

The challenge: Bobby is assisting his customers in facilitating payments for their customers. This is an immediate escalation in compliance requirements, because Bobby is inserting his business into the payment process, which contains sensitive, protected data (see Figure 4). Furthermore, Bobby is handling payments for multiple merchants, which means he must also be segregating and protecting data.

The solution: During the development of Bobby's application and infrastructure, it is crucial that data segregation and encryption are in place. This approach will help him adhere to necessary PCI standards, as well as others concerning data integrity and confidentiality. Furthermore, when Bobby's customers come to request their data, he must ensure that no commingling is happening.




Figure 4. Customer PCI data within the organization must be defended.

Unnecessary Data Exposure

The problem: Another issue that many organizations tend to gloss over is just how vulnerable they may be internally. To make life easier, when Bobby hired and set up accounts for his developers, he simply gave them all a shared administrative account (see Figure 5). Unfortunately, this is a dangerous practice that may result in security and/or data concerns.

Let's examine a few possibilities:

- The development team is likely working on the business during all hours of the day—and potentially on multiple devices. Bobby needs to gain insight into whether the data is being synchronized and/or used by his development team outside of his protected space.

- The sharing of credentials among the development team also poses a significant risk. For example, what happens when developers leave? Are their credentials changed?

The fix: Identity access management is one of the cornerstones of cloud infrastructure. For this reason, Bobby should take advantage of the robust authentication mechanisms put in place and ensure that his team uses unique credentials. Furthermore, he needs to ensure that users have only the privileges required.




Figure 5. The development environment has potential data exposure.

Defaults Don't Always Help

The Problem: One of the greater areas of risk that incident responders encounter as organizations deploy applications and solutions within the cloud is a reliance on technology defaults.

Unfortunately, many applications are designed to be open sourced, hacked together and then secured by the organization itself. Many NoSQL solutions, for example, used to be available with ports open and available to the internet (see Figure 6).

In early 2017, this led to a massive global issue of data compromise and NoSQL databases being held for ransom.

In Bobby's Bits, Bobby may not have hired the correct security personnel to help harden the various applications. Furthermore, if Bobby's development team simply was working on a fix, it may have

inadvertently left default ports and/or default accounts open and accessible to the world.

The Fix: The fix is twofold. First, Bobby and the development team should work to ensure that the various applications and tools they use are hardened by—you guessed it— compliance standards. This may include enabling encryption, setting up role-based access controls and limiting open ports/network routes to the application.

Additionally, with Bobby's infrastructure being in the cloud, he can resort to automated compliance scanning and verification tools. These scanning and vulnerability assessment tools will be kept up to date by the various vendors and can help ensure that the application is protected against the latest as well as pre-existing threats.

Furthermore, because Bobby's infrastructure is in the cloud, he can schedule scans much more frequently than some organizations like to do at the physical level.




Figure 6. The use of default settings can put the organization at risk.

Summary

Many decisions regarding data security and compliance are made utilizing standards set forth to help protect that particular content. Unfortunately, as organizations experience growth and network expansion, they often make decisions that may impact the safety and integrity of the data they store and use.

This is not an intentional mistake. Many organizations are growing exponentially and are seeking technology to facilitate that growth. This has driven a lot of organizations to the cloud—all in all, a great thing! The cloud can solve unique challenges of scale and availability—something very crucial to business. However, some organizations are also thinking that because the data is in the cloud, security is no longer their problem.

In this paper, we examined the concept of compliance-forward thinking, which asks organizations to consider compliance requirements when they are planning and building infrastructure, instead of afterward. There is a wealth of options within the cloud service space that can assist in automating and monitoring compliance of your organization and/or your customers' data.

As more organizations consider the options that cloud services can bring their business, it is crucial that compliance is at the top of the list of requirements. We have found that by starting the conversation with compliance in mind, what was once a tricky subject has become a guiding light to help organizations make safer decisions about the handling of customer data.

A few parting thoughts for organizations that are currently facing these issues head on:

- Look for areas within your cloud providers where compliance can be automatically monitored and/or reported on. Furthermore, look for compliant-ready deployments that can help fix requirements head on.
- Almost all compliance requirements include basic access rights monitoring, to ensure that employees are not sharing accounts and/or access mechanisms. If you set up individual accounts from the start, this requirement will already be fulfilled.

- It is easy to take newer technologies, drop them in place and begin working. But time and time again, we see organizations suffer breaches and noncompliance because of following the defaults. Make sure your team knows how to harden—and maintain—a good state of security within your applications and associated software.

About the Author

Matt Bromiley is a SANS Digital Forensics and Incident Response instructor, teaching Advanced Digital Forensics, Incident Response, and Threat Hunting (FOR508) and Advanced Network Forensics: Threat Hunting, Analysis, and Incident Response (FOR572), and a GIAC Advisory Board member. He is also a principal incident response consultant at a major incident response and forensic analysis company, combining experience in digital forensics, incident response/triage and log analytics. His skills include disk, database, memory and network forensics, as well as network security monitoring. Matt has worked with clients of all types and sizes, from multinational conglomerates to small, regional shops. He is passionate about learning, teaching and working on open source tools.

Chapter 6: How to Build a Data Security Strategy in AWS



Dave Shackleford

SANS Senior Instructor & Author

"It's clear that more organizations are moving sensitive data into the cloud, but what does this mean for us? Security professionals have enjoyed a wide range of security controls for protecting data on premises, including encryption, access controls, data loss prevention (DLP), classification and life cycle tools, and more. For quite some time, many of these controls weren't readily available in the cloud, but that time has passed.

Today, security teams have a great selection of security tools and controls for all different types of cloud storage and data usage services, as well as lots of ways to monitor data access and use. This chapter outlines the types of controls teams should consider for all aspects of data security in AWS."

The Importance of Data Security in the Cloud

Global organizations are adopting cloud solutions for a variety of compelling reasons, ranging from new business opportunities to reduction in costs to overall improvements in operational efficiency. That makes security in the cloud more important than ever.

In the Cloud Security Alliance's Top Threats to Cloud Computing research from August 2018, organizations ranked data breaches as the top concern for cloud deployments—no different from the major concerns for on-premises assets.¹ Naturally, this also means that as part of the shared responsibility model, organizations have the authority to enable controls in the cloud to protect data from exposure and attack.

The good news is that more data security controls and products/services are available than ever, and they are more fully mature.

In this paper, we break down key controls and considerations for protecting your data in the AWS cloud, including encryption and key management, data loss prevention, classifying and tracking data, and more.

The Kinds of Data We're Putting in the Cloud

As organizations put more sensitive data into the cloud, they are increasingly willing to better accommodate critical business needs by allowing such data in public cloud environments. In the most recent SANS cloud security survey, respondents from a variety of organizations worldwide indicated that they were storing business intelligence data (48%), intellectual property (48%), customer personal data (43%) and financial business records (42%), among many other types of data, in cloud environments.²

At the same time, organizations have a need to meet regulations and compliance requirements focused on data security. The same cloud security survey also revealed that, for more than half of respondents (54%), privacy regulations such as the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA) have impacted existing or planned cloud strategy, with another 12% unsure of impact.

When storing sensitive personal information in the cloud, it is imperative to choose a provider that can facilitate compliance to privacy regulations and has a global presence in the various regions needed

¹ "Top Threats to Cloud Computing: Deep Dive," <https://cloudsecurityalliance.org/artifacts/top-threats-to-cloud-computing-deep-dive> [Registration required.]

² "SANS 2019 Cloud Security Survey," May 2019, www.sans.org/webcasts/state-cloud-security-results-2019-cloud-security-survey-109760

to support these important regulatory requirements. Over time, it's likely that more and more region-specific privacy laws and requirements will come about, which will necessitate choosing cloud provider partners that can keep pace with these changing controls and reporting needs.

“As part of the shared responsibility model, organizations have the authority to enable controls in the cloud to protect data from exposure and attack. The good news is that more data security controls and products/services are available than ever.”

Critical Aspects of Data Security in the Cloud

Mature organizations today need to address many considerations to adequately protect data, and that applies for their cloud deployments. In the cloud, these considerations range from classification to implementation of various controls to governance and process adaptation within cloud engineering and operations teams.

Data Classification Policies

Identifying standard definitions for data is easy. Putting them into practice and maintaining them is never as simple, but tools are definitely emerging to classify and track data in the cloud. Amazon Macie is a security service that uses machine learning to automatically discover, classify and protect sensitive data in the AWS cloud.³ Amazon Macie can recognize sensitive data patterns such as personally identifiable information (PII) or intellectual property, and provides organizations with dashboards and alerting tools that provide visibility and insight into how this data is being accessed or moved. The service automatically and continuously monitors data access activity for anomalies based on usage profiles (both from individual accounts and metadata from the overall usage patterns of many accounts over time) and generates detailed alerts when potentially illicit access or data leaks are occurring.

³ This paper mentions product names to provide real-life examples of how security tools can be used. The use of these examples is not an endorsement of any product.

“When storing sensitive personal information in the cloud, it is imperative to choose a provider that can facilitate compliance to privacy regulations and has a global presence in the various regions needed to support these important regulatory requirements.”

Any organization planning to store sensitive data in the AWS cloud should strongly consider enabling Amazon Macie to profile and monitor data of specific classification types, and send Macie events to Amazon CloudWatch for even more detailed alerting and automation workflow enablement. And Amazon Macie data, like several other security services' output, can be sent to a new Amazon service called AWS Security Hub, which can aggregate security details across accounts and report on current security posture in a centralized console.

Types of Controls

Let's explore some of the types of controls and focal areas most organizations rely on today for data security in AWS.

Encryption

Encryption is a major area of interest for cloud implementations, primarily because it offers one of the few true lines of defense when moving resources into outsourced environments. All types of data encryption are encompassed, ranging from data at rest to data in motion and even data in use within applications. Some challenges come along with this, however.

For data at rest in the cloud, organizations have several major types of encryption to consider:

- **File/folder encryption** — File and folder encryption relies on applying a policy that dictates what to encrypt and who can access it.

- **Full-disk encryption for cloud workload storage volumes** — Full-disk encryption can help solve the problem of data exposure within virtual machines, but key management is a major concern.
- **Specialized encryption (database, email)** — Specific encryption for database columns or tables, as well as email stores, can be implemented in the cloud too.
- **Cloud-native storage encryption** — For specialized storage options like Amazon S3 buckets, encryption is easiest to implement through built-in AWS configuration options that allow for selection of encryption keys and access controls.

Each method has its pros and cons, and products and services are available in every category to assist in building a data encryption model that is sustainable and meets all necessary requirements. File and folder encryption products are generally compatible with cloud environments. For example, users with the appropriate rights to perform the encryption operation could easily encrypt files and folders in either a platform-as-a-service (PaaS) or infrastructure-as-a-service (IaaS) implementation. The encryption product would need to be present within the instance, however, and the user profile would need to retain some sort of key accessibility. This can be an issue for PaaS environments in particular, because user and role management systems may rely on vendor-specific APIs or internal systems that do not support the needed encryption key access. This can also be challenging for environments with numerous access types, such as partners, vendors and various internal roles.

For most organizations, enabling full-disk volume encryption for workloads in PaaS and IaaS implementations is an easy and relatively low-cost option. While not all of these encryption types truly support master boot record (MBR) encryption or granular recovery options, they really are not intended for this anyway (because these options are usually for mobile devices that could be lost). Instead, volume encryption protects any snapshots or replicas/backups taken automatically, and key management and integration are usually vastly simplified within the native cloud provider environment. In AWS, enabling Amazon Elastic Block Store (EBS) encryption is simple, using either the Amazon EBS customer master keys for the account or unique keys that are either uploaded into the AWS Key Management Service (KMS) or created there by the organization. Implementing the encryption is possible as a default option for all new workloads and storage volumes, or security teams can enable encryption on a volume in the web console in just a few steps.

Protecting data in motion is important for the cloud, primarily in two places:

- **Between the on-premises environment and AWS**, where sensitive data may be passing constantly in the case of hybrid architectures or intermittently for other cloud deployments.
- **Internally within the AWS infrastructure**, which would then rely on point-to-point tunnels between workloads, data encryption or both.

Amazon makes site-to-site encryption simple with IPsec VPN connectivity to a virtual private gateway (VPG) object within a customer's virtual private cloud (VPC). For more elaborate infrastructures, especially those with high-speed requirements or multiple inter- and intra-cloud connections, organizations may need customized hardware platforms and even acceleration solutions (available from a number of third-party vendors). Organizations can establish a true point-to-point private connection with the AWS Direct Connect service, too. This service provides a dedicated, guaranteed throughput connection to an on-premises environment, which functionally allows the AWS cloud to become an extension of the organization's network. One important point is that dedicated point-to-point services for network connectivity, such as AWS Direct Connect, are not natively encrypted—this is a common misconception! To encrypt data for transit across AWS Direct Connect links, organizations need to enable VPN tunnels within them, or perform application- or data-level encryption.

Managing, storing and controlling encryption keys are critical factors when using encryption in the cloud. AWS KMS is a managed hardware security module (HSM) service within AWS. It is possible to create keys in a region or import them from in-house key-generation solutions. Numerous AWS services are integrated with AWS KMS, including EC2 and S3. In fact, all major storage types within AWS now support various forms of encryption, all of which can be integrated directly with AWS KMS. Amazon's KMS also includes an in-depth audit trail with AWS CloudTrail, where all API requests and actions related to AWS KMS and key access are logged securely.

“One important point is that dedicated point-to-point services for network connectivity, such as AWS Direct Connect, are not natively encrypted— this is a common misconception!”

Amazon also has independent management and auditing within AWS, so there is strong and documented separation of duties in place within the environment. Numerous compliance certifications/assertions are also in place for AWS KMS. For customers that need even more control over keys, AWS CloudHSM is a full HSM that the customer can provision, enabling it to generate and use its encryption keys on a FIPS 140-2 Level 3-validated hardware platform. AWS CloudHSM protects your keys with single-tenant access to tamper-resistant HSM instances in your own VPC. You can configure AWS KMS to use your AWS CloudHSM cluster as a custom key store rather than the default AWS KMS key store, too, integrating the two services for simpler provisioning and use of keys within AWS storage services.

Data Loss Prevention

Data loss prevention (DLP) has been challenging for many organizations to implement in the cloud, primarily because of a lack of solutions and difficulty integrating with the cloud provider's APIs. That has significantly changed in the past several years, however. In addition to tools like Amazon Macie as a cloud-native option, quite a few third-party providers have added products and services in the AWS Marketplace to offer network DLP (usually through the implementation of a virtual gateway appliance), as well as host-based DLP agents that can be installed into workloads and images, reporting back to a central monitoring and policy platform also deployed in the cloud environment.

Implementing DLP is a subjective decision depending on whether your organization is subject to internal or compliance-related requirements that may necessitate this particular control, but there are products and services that can help you accomplish this if needed.

Data Life Cycle Controls

The most common data life cycle model has seven phases, as shown in Figure 1.

Generation

Phase 1 of the data life cycle is data generation. With regard to data generation and instantiation, security teams should focus on the following areas:

- **Ownership** — Who owns and maintains the data that moves to the cloud? This will likely be a business unit or some sort of cooperative effort between business and IT. Data owners have a bad habit of forgetting that they are the data owners (placing this burden on the data custodians), so it's a good idea to ensure that the actual stakeholders understand the risks and that they sign off on the level of cloud deployment and security controls needed to ensure the data remains safe.

- **Classification** — What types of data are we tasked with managing? Look at data classification policies and cloud-enabled tools and services to help track and monitor specific data types.
- **Governance** — Who is responsible for the data throughout the entire life cycle? Again, this could be one group or, more likely, a cooperative effort. For security professionals, ensuring data security throughout the entire life cycle (not just when it's generated) is a top concern

Use

Data use, the second major phase of the life cycle, involves the following major security concerns:

- **Data access** — Enable data access controls that align with least-privilege business use cases.
- **Legal access** — Determine whether the data will be accessible to legal counsel (for electronic discovery, for example).

It's a good idea when planning cloud deployments to build a map or breakdown of the data types that will be accessed and used in the cloud, where they will be stored and who will need access to them. This exercise also enables teams to do a much more effective job of creating role and privilege strategies.




Figure 1. Data Life Cycle Model

Transfer

The third phase of the data life cycle, data transfer, encompasses the movement of data between systems and applications. The fundamental concerns for this phase include:

- **Public/private networks** — What kinds of networks are involved in data transfer (public or private)? For a cloud implementation, a hybrid of internal and external network resources is likely. Anything going across the internet, of course, is a public network.
- **Encryption** — Is the data encrypted during transfer? Data can be encrypted before transit, sent through an IPsec VPN tunnel or both.

There are many options to control and encrypt data in transit, whether through using native cloud technologies or third-party tools and vendor products. Many firewalls can now be used to create and terminate VPN tunnels easily, too, so a cloud firewall strategy may be another possibility to help with this.

Transformation

Data transformation, the fourth stage, is where some sort of processing occurs, typically through the interaction with applications. The following are concerns and considerations during this phase:

- **Integrity** — How will data integrity be maintained in the cloud environment? Data integrity will be handled through SLAs to ensure no corruption or data loss occurs.
- **Sensitivity** — Will the data still be considered PII after modification? This classification largely depends on how the data is being sent to the cloud and processed. At one stage, it may be considered sensitive data, whereas at another it may be obfuscated or not have any recognizable qualities as personal or sensitive data.
- **Attribution** — Will the data be attributable to an individual or organization after transformation? Again, this will depend on the applications in use and the manner of storage.

Storage

Cloud storage (stage 5) is a concern for obvious reasons. We have covered encryption for data at rest, and this is one way to potentially offset some of the risks of sensitive data stored in a cloud environment.

Along with encryption and access controls, it's a good idea to check on the SLAs for resilience, availability and processing/transfer, as well as ensure you can export data easily as needed.

Archival

How is data backed up and archived? What are your data retention requirements for compliance and internal policy? For cloud implementations, consider the following areas during the data archival phase (stage 6):

- **Legal/compliance concerns** — How long must the consumer store the data? For example, log files for PCI DSS compliance must be retained for a year.
- **Storage types** — Different types of storage within AWS may be more suitable for longer-term archival. Amazon Glacier, for example, is an affordable way to perform backups and archive data in the cloud, but performance is more limited. The service has several security measures built in, including IAM-controlled access, automatic AES-256 encryption and TLS-encrypted endpoints for secure transfer (both from the internet and within EC2 workloads).

Destruction

The last major phase of the life cycle is data destruction. For the cloud, you need to think about:

- Getting a certificate of destruction from your cloud provider, if available
- Simply encrypting all of your data and then shredding the key as a means of ensuring the data is unrecoverable

Data can be recovered from AWS physically, too, by using the Amazon Snowball or Amazon Snowmobile service. Amazon Snowball is a petabyte-scale data transport solution that uses devices designed to be secure to transfer large amounts of data into and out of the AWS cloud. Amazon Snowball devices use tamper-resistant enclosures, 256-bit encryption and an industry-standard Trusted Platform Module (TPM) designed to ensure both security and full chain of custody for data, with all encryption keys stored in AWS KMS. The Amazon Snowmobile service is similar, but it is an exabyte-scale data transfer service used to move extremely large amounts of data to and from AWS via a 45-foot-long, ruggedized shipping container, pulled by a semi-trailer truck.

User Behavior Analytics + User Activity Monitoring

While not specifically a data security control, the need to monitor user access to data has grown exponentially in recent years as a result of account compromise, insider threats and many other attack vectors, all of which necessitate keeping a closer watch on data altogether. Within AWS, enable Amazon GuardDuty to monitor for unusual activity or behavior related to users and workloads. Amazon GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behavior to protect customers' AWS accounts and workloads. Amazon GuardDuty analyzes billions of events across multiple AWS data sources, such as AWS CloudTrail, Amazon VPC Flow Logs and DNS logs.

Differences in Security Controls for Hybrid Architectures

A number of data security concepts change in a hybrid architectures model. Some of the following are the most important to consider when building and planning your cloud architecture and operations strategy:

- **Cloud provider SLAs and data availability/resiliency guarantees are now a part of your shared responsibility strategy.** For example, many AWS S3 and S3 Glacier storage types offer 99.999999999% durability of objects over a given year (that's 11 nines). Most uptime guarantees are 99.5% and above, and service credits may be contractually guaranteed when these are not met. (Be sure to discuss with AWS beforehand and understand all contract terms.) This is a prime example of shifting some of the traditional responsibility of service uptime and integrity to the cloud provider. Being able to share the risk by transferring to the provider some (not all) responsibility for data availability and resiliency can possibly free some operational capacity to implement and maintain additional data security controls.
- **Secure transport of data is critical across certain data paths.** While secure transport of data has always been important, creating a hybrid architecture requires transport of data across the internet, an untrusted network. Fortunately, between dedicated connections like AWS DirectConnect and industry-standard site-to-site encryption with IPSec, secure transfer of data is easy to accomplish in a hybrid architecture. Using third-party encryption gateways or network gateways can also facilitate secure data transfer in a larger deployment.
- **Use of cloud-native data security controls is likely a requirement.** Plenty of data security options are available in the AWS cloud, both from AWS and third parties. However, at least some of the cloud-native controls, such as AWS KMS, are likely needed to facilitate

implementation of encryption easily. Other cloud-native services related to data security may be more affordable and easier to implement in AWS. These include AWS Certificate Manager (ACM) for the creation and life cycle management of digital certificates and AWS Secrets Manager for secure storage of keys and credentials used in provisioning system and data access in workloads, DevOps pipelines and more.

- **Emphasis on bring your own key (BYOK) and better encryption oversight will be paramount.** Today, AWS readily supports import of keys generated on your own premises, which may be a regulatory requirement or internal best practice. Having industry-leading encryption storage available through HSMs may also facilitate better audit controls for keys and key access, as well as key life cycle management. Given the increasing use of encryption as a core data security control in the cloud, flexibility in key generation, storage and life cycle management are need-to-have requirements for more organizations today.
- **Technology needs to work internally and in the cloud in some cases.** When using a hybrid architecture, you will already have some data security controls in place in your internal environment, and for a variety of reasons, you may need or desire to continue using products and services from third-party providers. Fortunately, an increasing number of providers have partnered with AWS through the Marketplace program to offer data security controls that can natively work in AWS alongside your existing implementations.

While some of these changes and shifts will be harder to accomplish than others, all are important to consider when building a hybrid architecture.

Scaling Your Data Security Strategy to the Cloud

When moving to the cloud, or expanding your footprint within AWS, it's important to know your data and look at tools and tactics to track your data in the cloud. Even if you don't need full-fledged DLP tools (which are available), monitoring and tracking specific data types and access to these data stores can significantly enhance your data security and privacy strategy altogether. Tools like Amazon Macie can enable this capability for your organization simply and effectively, and you can then build specific monitoring workflows for alerts from this service to detect illicit access or patterns of access that may indicate insider abuse or compromise.

Implementing encryption in and to the cloud for transport and storage is a requirement for most organizations today, and the use of encryption will only continue to grow. The earlier you plan to leverage in-cloud tools and services to enable encryption (key creation, storage, access, auditing and

life cycle management), the more empowered you will be as your cloud deployment expands. AWS KMS, for example, is integrated with all AWS storage models and can be used to store, create, audit and destroy keys. AWS CloudHSM provides an additional layer of security with dedicated hardware that also integrates with AWS KMS if needed. By updating your key creation, import and life cycle policies and processes to incorporate these cloud-native technologies where appropriate, you will be far better prepared to expand encryption use as needed.

Ensure you have access controls on data stores and monitoring through audit logs, because all sensitive data access within the cloud environment should be monitored and controlled. Many of the storage types in AWS have access controls that can be enabled, and all data and storage access can be monitored through AWS CloudTrail. Amazon S3, for example, has the following controls related to access control and auditing.

Data access:

- IAM policies — User-, group- and role-based access control to storage buckets
 - Bucket policies — Policies applied to a specific S3 bucket and nowhere else
 - ACLs — Bucket- and data-specific access controls for users/groups
 - Query string authentication — REST-based access key strings that can be passed to AWS for access control
- **Access logs:** All S3 access and activities can be logged to a separate bucket for collection and analysis.

“The earlier you plan to leverage in-cloud tools and services to enable encryption (key creation, storage, access, auditing and life cycle management), the more empowered you will be as your cloud deployment expands.”

Two new features added to Amazon S3 in 2018 are critically important and can enhance S3 deployments' security posture enormously. First, S3 Block Public Access is a default deny model for an entire account that organizations can turn on to prohibit any S3 bucket from being made public. Amazon S3 Object Lock can turn an S3 bucket into a write-once, ready-many (WORM) system, useful for legal retention of data and evidence in chain-of-custody cases, too.

As another example, the Amazon Relational Database Service (RDS) offers the following access security features:

- **DB Security Groups** — Similar to AWS EC2/VPC Security Groups, these are network ingress controls that you can enable by authorizing either IP ranges or existing Security Groups. These allow access only to the database port(s) needed and do not require a restart of the database instances running.
- **IAM permissions** — Can be used to control which Amazon RDS operations each user can call.

Security teams should enable a least-privilege access model for all storage services used within the AWS cloud, and also make sure to turn on AWS CloudTrail and any additional logging.

Finally, plan for all phases of the data life cycle, from creation through destruction, as well as changes to how data may be handled and controlled over time. In the cloud, there are many more storage and data control options than you likely have accessible in-house, and you can leverage a hybrid data life cycle strategy across them. For example, an organization may store certain sensitive data in Amazon S3 for a year to meet PCI DSS access requirements, but then move the data to Amazon S3 Glacier after a year to save money (where access is slower, but no longer required for compliance).

Case Study: Data Security Operations in a Hybrid Architecture

Acme Corp. was planning a significant cloud migration to AWS and wanted to ensure that it didn't skip or fail to implement any important data security controls and processes that could negatively impact compliance. Additionally, Acme viewed a move into AWS as an opportunity to review data security controls and practices at the corporation and hoped to improve its security posture in many ways by taking advantage of many cloud-native options.

First, Acme reviewed its existing data security and data classification policies to ensure that the language in place accommodated cloud use cases. It determined that it was comfortable moving all but its most critically sensitive data to the cloud to start and that it could revisit this decision periodically after it had things up and running smoothly. Personal data on customers would be migrated, as would some business financial data and human resources databases.

To prepare for data security in the AWS environment, the team enabled a BYOK strategy using AWS KMS. Within AWS KMS, Acme chose a default expiration date for keys of six months to start—AWS KMS even generated an automatic Amazon CloudWatch metric that tracks each key's expiration to alert Acme! The enterprise security operations team that maintains the internal HSM at Acme updated its rotation and key management processes to incorporate the use of AWS KMS, with console and AWS Command Line Interface (CLI) operations documented to create new keys, upload them into AWS and monitor for key life cycle thereafter. The team determined that it did not need to use AWS CloudHSM at the moment, but it decided to revisit that later as well, especially if/when Acme opted to move its most sensitive data into AWS.

For compliance and internal requirements, the team decided that it needed to implement a DLP solution in AWS. Acme's existing in-house provider is an industry leader in the space, and the team preferred to continue using this solution if possible. After investigating options, it found that the third-party solution was available in the AWS Marketplace, and Acme would simply need to license a new virtual image deployed in the cloud.

To take advantage of many of the security features in AWS, the team selected Amazon S3 as the main storage location for some of the most sensitive data, primarily to take advantage of Amazon Macie for monitoring and reporting on sensitive data access.

The S3 Block Public Access policy was enabled for Acme's account, and specific access controls were created to enable a least-privilege access model through IAM privileges. Amazon S3 bucket logging was also enabled, and AWS CloudTrail was turned on to further monitor all access to assets in the VPC. The team also enabled Amazon GuardDuty to track account activity and behavior as the number of users and groups using AWS grows.

For all EC2 instances, the team enabled default Amazon EBS volume encryption using AWS KMS keys that it had uploaded from Acme. For all RDS databases, column-level encryption was implemented where needed, and Security Groups controlled network access to the databases as well.

All AWS VPC connectivity needed to be secured as well, because Acme chose to implement a hybrid architecture. The team easily accomplished this by setting up an IPSec tunnel between Acme's on-premises network gateway and the VPG within the VPC. As the environment grows, it's likely that Acme will implement a DirectConnect pipeline, too, but this will come in the next deployment phase.

Summary

Securing data in the cloud is easier than ever, largely because of the plethora of cloud-native controls and tools available. For many organizations, it's just a matter of choosing the right combination of controls and services to meet their business and operating requirements. Encryption, access control and monitoring are all available readily within the AWS cloud. Encryption key storage and life cycle management are easily managed, but they require planning and likely adapting existing processes to use in-cloud platforms like AWS KMS and AWS CloudHSM. Tracking sensitive data access is possible at scale with services like Amazon Macie, and monitoring all user behaviors (for data access and more) is easily done with Amazon GuardDuty. Protecting data at rest, in transit and in use has always been, and will continue to be, a major priority for security teams. In the AWS Cloud, there are numerous ways to accomplish this.

About the Author

Dave Shackelford, a SANS analyst, senior instructor, course author, GIAC technical director and member of the board of directors for the SANS Technology Institute, is the founder and principal consultant with Voodoo Security. He has consulted with hundreds of organizations in the areas of security, regulatory compliance, and network architecture and engineering. A VMware vExpert, Dave has extensive experience designing and configuring secure virtualized infrastructures. He previously worked as chief security officer for Configuresoft and CTO for the Center for Internet Security. Dave currently helps lead the Atlanta chapter of the Cloud Security Alliance.

Chapter 7: How to Design a Least Privilege Architecture in AWS



Dave Shackleford

SANS Senior Instructor & Author

"Identity and access management (IAM) is one of the lynchpins of a sound cloud security strategy, but many organizations struggle with this complex topic. Creating a sound set of roles and identity policies for both user and service access in AWS is a critical area of focus for security teams, alongside network access control with cloud-native microsegmentation. Rounding out the core pillars of least privilege architecture for AWS is cloud security posture management, helping to ensure the control plane itself is defended.

This chapter breaks down all these areas and more, enabling security teams to enable least privilege access models throughout the entire cloud."

Introduction: What Is Least Privilege?

On the surface, the concept of least privilege is seemingly obvious: In any given scenario or use case, only allow a user, service, application or system to operate with the bare minimum privilege necessary to successfully accomplish the business goals desired.

However, over decades of computing, consistently implementing least privilege as a best practice has been a challenge for a number of reasons, including:

- **The ability to determine the appropriate “least privilege” for a given use case is a surprisingly complex issue.** It's often challenging for administrators, engineers and developers to plan for and think through the exact set of privileges needed to implement a least privilege access model because of widely differing access needs from different types of users and services.
- **It is easier to allocate more privileges than to limit access.** Security professionals have observed this classic problem in many different scenarios over the years, ranging from data center administration to development and application interactions to end users on their workstations. It's much more convenient for workers to do whatever they need to when they are assigned extensive privileges.
- **The range of permissions and privilege models varies widely between environments and applications/services.** Because there is little to no commonality across the use cases and technologies we employ from one organization and environment to the next, developing a consistent model of least privilege can be time-consuming.

That said, even successful least privilege implementations tend to shift and drift over time without continuous monitoring and oversight.

Least Privilege Concepts in the Cloud

Security professionals are rethinking the approach to least privilege security concepts for the public cloud. Some key factors to address include:

- **Vanishing perimeter** — The cloud is a cohesive ecosystem that relies on numerous service and application interactions, and the classic idea of the perimeter is changing.

- **Application workloads** — Security professionals need a better understanding of application behavior at the workload level. They should be looking at the types of network communication approved applications really should be transmitting.
- **Trust relationships** — The focus should be on trust relationships, system-to-system relationships and service-to-service relationships within all parts of the cloud environment. Most communications in enterprise networks today are either wholly unnecessary or irrelevant to the systems.

Pillars of Least Privilege

Security teams need strong access controls to effectively secure who can do what and from where. The “who” could be a user or app identity or systems/ subnets within the environment. Many cloud access management strategies are starting to revolve around the idea of least privilege at all layers, which some may call “microsegmentation” or a “zero trust” design. Whatever you choose to call it, the three elements of this strategy, illustrated in Figure 1, are:

- Identity and access management (IAM)
- Network access and segmentation design
- Cloud security posture management



Figure 1. Least Privilege Pillars

“Security teams need strong access controls to effectively secure who can do what and from where.”

For many organizations, designing least privilege access controls often encompasses a blend of cloud-native and third-party controls as well. This area is evolving quickly, so security teams should pay careful attention to the market and open source communities too.

Although the first two pillars are more critical, all three are needed to enable a comprehensive least privilege strategy. In the coming pages, we explore the three pillars of the least privilege model and look at how they can work together to implement an effective least privilege strategy.

Least Privilege Pillar 1: Identity and Access Management

Arguably, one of the most important aspects of cloud security is IAM. If you think about it, IAM is a linchpin to controlling most elements of security for who and what can access resources in the cloud. Defining roles, enabling strict access models and limiting the resources available to users and systems is a critical step in enabling a sound cloud security strategy overall. A key element of IAM that security teams need to adapt to is the use of IAM for enveloping assets, allowing them to create least privilege architectures with affinity policies in place.

User Relationships

IAM users are associated with credentials for making API calls to interact with cloud services and exist only within the cloud environment itself. By linking directory services like Active Directory to the cloud, security teams can leverage existing in-house users and map them to IAM groups and roles, but a standalone user created within the cloud is only useful in the cloud. New IAM users have no permissions (an implicit “Deny All” policy). This is a good thing, because permissions must be explicitly granted. This policy can also help with the common problem of over-allocating privileges to users and groups in the environment.

IAM users can represent any asset/resource—an IAM user is a simple identity with associated permissions. This means that IAM users can be enabled for application access to Amazon Web Services (AWS)¹ resources too, not just as actual interactive user accounts. Once you create service-oriented users, place them in defined groups, if warranted. Security teams can assign permissions and privileges directly to users (not advised) or groups (better to manage and maintain).

Service Relationships

For service interactions within the environment, however, cloud security teams should focus on defining specific roles. There are four types of roles:

- **AWS services** — This type is for provisioning roles that will be assigned to AWS services like Amazon EC2 and AWS CloudFormation. In other words, what resources can access other resources in AWS, and what actions can they take? This type of role forms the basis for instance profiles, which we cover in a moment.
- **Cross-account access** — Teams can provision access to their AWS infrastructure to other AWS accounts the organization owns or to third-party AWS accounts.
- **Federation** — For federating access with SAML 2.0 to in-house directories, a federation role is available.
- **Identity providers** — These role types work with identity providers (IdPs) for single sign-on (SSO) and federated access to resources. There are three types of IdP roles. The first focuses on web IdPs like Google, Facebook and Amazon Cognito. The second grants web-based SSO to Security Assertion Markup Language (SAML) providers, likely some of the most common for management console access. For direct SSO access to APIs via SAML, a third type of IdP role is available.

There are several distinct types of identity-focused least privilege orientation for cloud deployments and infrastructure. First, there should be a focus on any privileged users that need access to the cloud environment for administration, engineering or security-focused tasks. Ideally, even in large organizations, this should be a relatively small number of users that are carefully set up and monitored. The best practice for these users is to federate their internal user accounts directly to an assigned role within the cloud environment that has the fewest privileges assigned.

¹ This paper mentions solution names to provide real-life examples of how cloud security tools can be used. The use of these examples is not an endorsement of any solution.

The second major type of least privilege access model that all organizations need to consider is associated with deployment pipelines and associated systems and services. Whether on premises or fully hosted within the cloud environment, deployment pipelines need certain privileges to update workload images and containers, access code repositories, assign metadata tags to resources and monitor performance and security metrics and activities.

The third major type of least privilege focus is mapping user, service and application relationships wholly contained within the cloud environment. These might be Amazon EC2 workloads with instance profiles assigned that allow access to other AWS services like Amazon S3 buckets, AWS Lambda functions that need to interact with Amazon CloudWatch logs and database services, or service IAM accounts/groups used to allow access between applications and services in the environment.

Finally, privileges should be carefully reviewed for accounts accessing other accounts' services when a multi-account strategy is in place.

Relationship Mapping

For all of these different least privilege scenarios, organizations need to successfully map user and service relationships to create the most restrictive privilege models needed. Fortunately, a number of tools are available to accomplish this. During AWS IAM account creation, admins can use the AWS Access Advisor feature. Access Advisor shows AWS services allowed by the assigned IAM policy, policies assigned that grant specific permissions and last access times (if relevant). This information is especially helpful for users that are members of multiple groups with a variety of different policies in place. Many organizations have numerous groups, users and accounts that need to be handled differently, and it can get confusing. With this feature, admins can get a sense of what permissions are being applied, ideally before they are. The AWS Trusted Advisor service also informs account owners of some well-known privilege allocation issues that may be present.

“When using a multi-account strategy, review for accounts accessing other accounts' services.”

AWS IAM Access Analyzer, a feature within AWS Identity and Access Management (IAM), performs a more thorough analysis of privilege models in use. This tool helps organizations identify potential security risks in the AWS environment by analyzing the resource-based policies applied to resources within their zone of trust (the current account). When AWS IAM Access Analyzer identifies any policy that allows access to those resources by a principal that isn't within the zone of trust, the service generates a finding/alert. Security teams can use the information in each finding, such as the resource, access level and principal that has access, to determine whether the access is necessary or unintended. If the access is unintended, and therefore a risk, security teams can modify the policy to remove the access and work toward a least privilege identity model.

With an isolation and segmentation technique, each account is a completely isolated set of resources that can be configured to access resources in other accounts. For multi-account strategies employed to limit the post-compromise risk and provide highly granular least privilege access models, AWS IAM is a critical element of managing the access between accounts. AWS Organizations is a service that organizations can use to define policies and guardrails to apply across multiple AWS accounts from a master control level.

With AWS Organizations, you can create service control policies (SCPs) that really govern the use of other IAM policies. AWS Organizations can control the entire account, group and role life cycle with regard to policy application, and can do so for accounts that need to interact or have some relationship. Some basic examples of how AWS Organizations could be practical would be governing business unit (BU) account use (because they may have totally different needs, but still need some central control or billing), as well as governing and controlling DevOps and other team accounts (for the same reasons). AWS Organizations is the linchpin of a multi-account scope of impact limitation strategy in AWS—limiting the scope of impact to the smallest possible surface area prevents attackers from leveraging one compromised asset to access another. Creating a centralized policy model within AWS Organizations can allow security administrators to create different and least privilege policies for the appropriate accounts and assign them and/or revoke them easily. The service also provides a “master” rollup account that is often also the “payer” account that gets the consolidated billing for AWS accounts.

Setup and configuration of multi-account architectures have long been considered challenging and complicated tasks, especially for large organizations. Fortunately, numerous services and design models have been created within AWS to help with this. A sample multi-account framework to start from, called a “Landing Zone,” was proposed by cloud engineering experts several years ago, but creating and managing even this led AWS to create a new service, called AWS Control Tower, that can automatically deploy a multi-account starting architecture. Enterprises can then use AWS Control Tower to create and implement defensive guardrails such as AWS Config monitoring rules, infrastructure-as-code definitions

in AWS CloudFormation, and strict identity policies that restrict permissions and privileges across accounts, enable data encryption and much more.

Least Privilege Pillar 2: Network Segmentation for Access Control

The second major component of a traditional least privilege design model is network segmentation that is closely aligned with a specific type of system or workload. This is often termed “microsegmentation.” A least privilege concept of network segmentation strives to prevent would-be attackers from using unapproved network connections to compromise systems, move laterally from a compromised application or system, or perform any illicit network activity regardless of environment. By potentially eliminating lateral movement, a least privilege microsegmentation model also reduces the scope of impact when an attacker has illicitly gained access to an asset within a data center or cloud environment.

The classic model for implementing least privilege at the network level starts with a network access control policy of Deny All and then adds only those types of network access needed.

Microsegmentation with Cloud-Native Controls

The first category of focus for any cloud network isolation and segmentation should be the core network zone associated with cloud accounts. In AWS, this is known as the virtual private cloud (VPC), and this can contain any number of distinct network subnets. Cloud-native access controls can be created and applied within the VPC and should be used for isolating and controlling traffic flow into the VPC subnets altogether, as well as to and from instance workloads running applications and services.

“A least privilege concept of network segmentation strives to prevent would-be attackers from using unapproved network connections to compromise systems, move laterally from a compromised application or system, or perform any illicit network activity regardless of environment.”

Table 1. Differences Between Security Groups and NACLs	
Security Groups	NACLs
Apply to instances	Operate on VPC subnets
Only support Allow rules (layered on a default Deny)	Support both Allow and Deny rules
Are stateful	Are not stateful
Are considered in their entirety before traffic is allowed	Are processed in numerical order
Must be associated with an instance to apply	Apply automatically to all instances in a subnet

AWS has two built-in types of network access and isolation controls: security groups and network access control lists (NACLs). Use security groups and NACLs to control traffic into and out of network deployments. Security groups apply to instances and are stateful, whereas NACLs apply to VPC subnets and are stateless. Table 1 provides a breakdown of security groups versus NACLs.

It's a good idea when planning cloud deployments to build a map or breakdown of the data types that will be accessed and used in the cloud, where they will be stored and who will need access to them. This exercise also enables teams to do a much more effective job of creating role and privilege strategies.

In general, it's best to sparingly apply NACLs to either allow or deny known trusted or malicious IP addresses and subnets. The majority of the network access controls (NACs) should be defined and applied through the use of security groups. Because security groups begin in a "default Deny" state, it's much easier to create a least privilege model with them. Security personnel can enable Amazon VPC Flow Logs to track communications between assets in a VPC to ensure that no unusual or unexpected access is allowed, but a more complete coverage option to audit large numbers of security groups is to look into third-party network policy analysis tools that can ingest security group definitions and analyze them at scale.

Advanced Network Security Segmentation and Access Controls

To segment and control traffic at the application layer, or define policies focused more on application details and protocols, a third-party solution likely makes more sense, and many cloud options are available for enterprise-class networking. Most major cloud providers offer enterprise-class solutions that are capable of providing more granular policies and monitoring. Today's next-generation firewall (NGFW) platforms are often used to provide network intrusion detection and prevention, traffic inspection and behavioral monitoring, and centralized configuration and administration alongside existing on-premises NGFW platforms if desired. Leading providers include Palo Alto Networks, Fortinet, Sophos and others.

Segmentation/Isolation Best Practices

There are many well-known security fundamentals that organizations can follow when planning for and implementing least privilege network isolation and segmentation in the cloud.

First, be sure to consider what types of architectures make the most sense. For example, you can create all distinct assets in one very large VPC and control access with security groups and NACLs, or create a much more granular isolation strategy with multiple accounts and VPCs. Most major cloud providers support the concept of peering between virtual networking boundaries. VPC peering enables organizations to couple distinct VPCs together, allowing assets in one network to talk to assets in another. This capability can be incredibly useful in a design model because you can create true hub-and-spoke network designs that require traffic to pass through a transit zone of some type (through a dedicated security zone with intrusion detection and other controls, for example).

VPC peering is not transitive (i.e., there is no need to specifically allow it for each VPC peered together). In this case another type of platform, called a “transit gateway,” can simplify multi-VPC architectures significantly. This resource, which can be managed through the AWS Resource Access Manager service (for managing assets across accounts), can help teams create a more traditional hub-and-spoke model of network connectivity that will then have security groups and NACLs applied as needed. Much like route control, transit gateways can have IPS or firewall appliances attached as well, making these ideal for a central security control point. For managing multiple transit gateways, the AWS Transit Gateway Network Manager (AWS Network Manager) service enables organizations to manage all connected hybrid cloud network zones connected to and through transit gateways in a single dashboard. In many cases, teams set up a “transit VPC” with an NGFW platform as described earlier to process traffic to all other zones peered within the network architecture.

To summarize how IAM and core networking controls can facilitate a least privilege cloud deployment, be sure to:

- **Plan IAM roles and permissions to protect access to and use of VPC resources and services.** Many VPC objects and services can easily be controlled through IAM, including EC2 workloads, containers and much more.
- **Leverage security groups and NACLs to the full extent.** These controls provide built-in cloud-native NACLs to workloads and between subnets. If you need more control (and you likely will), consider a third-party virtual firewall/IPS appliance as a gateway.

Least Privilege Pillar 3: Cloud Security Posture Management

Cloud security posture management (CSPM) tools can assess the actual control plane of the cloud environments in use for compliance assessment, operational monitoring, DevOps integrations, risk identification and risk visualization. A CSPM platform should continuously monitor cloud security risk and potentially implement configuration changes in the cloud environment that facilitate least privilege access and much more. These tools also offer threat detection, logging and reports. In addition, they usually provide automation to address issues ranging from cloud service configurations to security settings as they relate to governance, compliance and security for cloud resources. Because many cloud platform settings relate to networking and IAM configuration, having a continuous monitoring engine that highlights over-allocation of privileges and permissive traffic policies can be invaluable.

Having interoperability between monitoring and automation is a critical advantage of a CSPM.

For enterprises grappling with hybrid architectures and container environments, where misconfiguration is a common threat to cloud security, a CSPM tool is an excellent step toward implementing continuous monitoring and alerting for the cloud provider fabric configuration. Common misconfigurations tend to be present with identity controls, workload security, logging enablement, network configurations and more. Organizations that are moving to or currently in hybrid deployment scenarios should strongly consider CSPM tools.

A Least Privilege Use Case

For an organization planning on deploying to a platform-as-a-service (PaaS) or infrastructure-as-a-service (IaaS) cloud environment with a focus on least privilege, there are multiple recommended steps:

1. Identify roles and responsibilities for team members requiring access to the cloud infrastructure.
2. Determine the type of network access needed.
3. Evaluate IAM roles and privilege assignments.
4. Monitor the cloud control plane.

In the first step, those responsible for the least privilege strategy carefully identify roles and responsibilities for any cloud engineering, DevOps and security team members that may need access to the cloud infrastructure. This should always be the first priority because the account owner (the root account) is one that should be almost wholly disabled, other than for billing and a potential “break glass” scenario if disaster strikes. These privileged users should be established through federation and role integration if possible, or standalone users within defined groups if not. It’s best to assign the cloud provider’s predefined policies that match these administrative roles whenever possible because these are likely to be the most accurate and well-structured. Even with that said, some of these can be used as beginning policies from which to reduce privileges as needed.

The next step is to determine what type of network access is required from the internet, from a hybrid cloud dedicated network connection and within the cloud environment itself. This step requires a review of application and service architecture to define data flows with TCP/UDP ports and application behavior profiles that planners can use to carefully restrict the types of traffic needed for operations. Organizations should plan to start with cloud-native networking controls like security groups and NACLs, which allow for a strong microsegmentation approach that can be managed through infrastructure-as-code (IaC) templates, such as AWS CloudFormation or HashiCorp Terraform, and monitored through API logs and metadata queries. For more robust network security, many enterprises will want to adopt a VPC peering arrangement for additional isolation, possibly with a third-party NGFW platform introduced to provide additional application-layer protection.

Throughout this entire process, identity, development and security teams should evaluate IAM roles and privilege assignments for workloads, services and all interaction between assets in the environment. Fortunately, tools such as AWS IAM Access Analyzer can be used to perform a deep dive into assigned roles and privileges for all components within a defined trust zone such as an account. Access logs: All S3 access and activities can be logged to a separate bucket for collection and analysis.

All teams involved should be invested in leveraging reports and alerts from tools like this to continuously look for opportunities to reduce privilege allocation wherever possible. This is an ongoing effort that will likely continue over time, because applications and assets continuously change and update within dynamic cloud environments.

Finally, it’s a good idea to consider a CSPM platform to continuously monitor the cloud control plane itself, looking for exposure and potential configuration pitfalls that could inadvertently allow for unintended or privileged access into services or the environment as a whole, as well as internal mappings of network and identity orientation that may be improved upon. For large, complex deployments, these types of third-party solutions can provide an extra set of eyes and ears on the cloud deployments overall.

Conclusion

A least privilege cloud architecture should include authentication and authorization controls, network access and inspection controls, and monitoring/enforcement controls for both the network and workloads. No single technology currently will provide a full least privilege design and implementation—organizations need to implement a combination of tools and services to provide the full degree of coverage needed. For most organizations, a hybrid approach of both cloud-native and third-party controls will make the most sense.

To implement a least privilege cloud environment, start with user and administrative access, followed by multi-account identity management, if applicable. From there, focus on network architecture and access control design, using cloud-native controls as the first line of defense and applying third-party controls for more robust defenses. Throughout all deployments, continuously evaluate privilege allocation and role assignments to find potential over-allocation of privileges where they may exist.

Once the cloud environment is up and running, a CSPM platform may make sense to continuously monitor the configuration.

More tools and services are available than ever before to aid in building and maintaining a cloud infrastructure that adheres to the principle of least privilege. A commitment to continuous oversight is critical because cloud environments tend to change rapidly. Implement tools as needed to provide adequate logging and alerting to ensure security teams are aware of how the environment is operating at all times.

About the Author

Dave Shackleford, a SANS analyst, senior instructor, course author, GIAC technical director and member of the board of directors for the SANS Technology Institute, is the founder and principal consultant with Voodoo Security. He has consulted with hundreds of organizations in the areas of security, regulatory compliance, and network architecture and engineering. A VMware vExpert, Dave has extensive experience designing and configuring secure virtualized infrastructures. He previously worked as chief security officer for Configuresoft and CTO for the Center for Internet Security. Dave currently helps lead the Atlanta chapter of the Cloud Security Alliance.

Chapter 8: How to Secure App Pipelines in AWS



Dave Shackleford

SANS Senior Instructor & Author

"This chapter focuses on the changing nature of application development and deployment—namely, more dynamic and more agile, with DevOps necessitating a heavy emphasis on monitoring, secrets management, privilege management, and many other security measures. For security teams to properly secure the CI/CD pipeline without creating barriers and roadblocks, a variety of governance practices are necessary, and more automated controls that integrate with development pipelines are important, too. In fact, the best app pipeline security controls and processes are embedded to create a DevSecOps model and culture."

We are seeing nothing less than an evolutionary shift as security infrastructure moves to software-defined models that improve speed and scale, and afford enterprise IT more agility and capabilities than ever before. Application development and deployment are driving this shift, and as the pace of development increases, organizations have a real need to ensure application security is embedded in all phases of the development and deployment life cycle, as well as in the cloud during operations.

Much like other areas of security, the responsibility for application security varies in the cloud widely, depending on the model in place. In a software-as-a-service (SaaS) model, the provider is entirely responsible for application security in almost every case. With a platform-as-a-service (PaaS) model, the provider supplies the underlying systems and templates, so it has a significant degree of control and responsibility—although any applications developed by the consumer are necessarily the consumer's own responsibility, and that extends to the security. With an infrastructure-as-a-service (IaaS) model, entire workloads and their contents (including application components) are the responsibility of the consumer.

In this paper, we delve into the changing nature of application development and security as organizations are building and deploying applications for the cloud. We'll cover the various phases of a modern application pipeline and discuss some of the security controls that organizations should consider implementing in each. We'll also touch on a number of other critical areas such as privilege management, containers and orchestration, and automation.

How the SDLC Is Changing

The software development life cycle (SDLC) has moved to a methodology that prioritizes collaboration and more frequent (yet smaller) updates to application stacks. Standards for code quality and security, as well as application workload configuration, should be defined and published so that all teams have something to measure throughout the entire application life cycle. Ideally, organizations will lock down cloud workloads as much as possible, running a minimum of necessary services. They should also revisit configuration requirements to ensure that any cloud-based infrastructure is resilient.

To shift toward a more collaborative culture, security teams need to integrate with the developers responsible for promoting code to cloud-based applications. Security teams can impress upon development and operations that they bring a series of tests and “quality conditions” to bear on any production code push without slowing the process. Security teams should work with quality assurance (QA) and development to define certain parameters and key qualifiers (such as bug count and severity) that need to be met before any code is promoted.

In addition, security teams need to determine which tools they can use to integrate into the application pipeline. They also need to identify areas and controls that may need to be updated or adapted to work in a Continuous Integration and/or Continuous Delivery model (covered in the next section). It is likely that new standards for many security prevention, detection and response capabilities should be revisited, as well. Examples of these areas include encryption, privileged user management, network security access controls, event management, logging policies and incident response strategy.

Once initial processes, policies and standards have been defined and agreed upon, the security team should focus on automation and seamless integration of controls and processes at all stages of the deployment pipeline.

The Modern CI/CD Pipeline

Many organizations are adopting Continuous Integration (CI) and Continuous Delivery (CD) for their cloud application pipelines. CI is often the most feasible part of the application development life cycle to be targeted by a team looking to speed up and implement more collaborative development practices. With CI, all developers have their code regularly integrated into a common mainline code base. This practice helps to prevent isolation of code with individual developers and can also lead to more effective control over code in a central repository.

CD is usually exhibited through small, incremental and frequent code pushes (often to stage or test environments), as opposed to the more traditional way of pushing code as large releases to production every few weeks or months. Modern development practices (e.g., Scrum, Kanban, Crystal, etc.) often release code more frequently than older models (e.g., waterfall) in an SDLC. CD means you deliver code to production in an automated pipeline, which is less common in traditional enterprises.

“The SDLC has moved to a methodology that prioritizes collaboration and more frequent (yet smaller) updates to application stacks.”

Modern cloud application pipelines strive for a number of goals and focal areas:

- **Automated provisioning** — The more automated the provisioning of resources and assets is, the more rapidly the SDLC and operations model can operate.
- **No-downtime deployments** — Because cloud services are based on service-oriented costing models, downtime is less acceptable.
- **Monitoring** — Constant monitoring and vigilance of code and operations help to streamline and improve quality immensely.
- **Rapid testing and updates** — The sooner code flaws can be detected, the less impact they'll have in a production environment. Rapid and almost constant testing needs to occur for this to happen.
- **Automated builds and testing** — More automation in the testing and QA processes will help to speed up all activities and improve delivery times.

Protection for application workloads requires a dedicated commitment to security at many levels of any organization. A sound governance model that includes collaborative discussions about code quality, system builds, architecture and network controls, identity and access management, and data security is critically important to developing the standards for controls and security posture (mentioned earlier).

Ideally, the following types of roles will be a part of any cloud application security and development model:

- Application development teams
- Cloud architecture and engineering teams
- Security architecture and operations teams
- IT in infrastructure teams (server engineering, database management and more)
- Compliance and legal teams (where appropriate)
- Business unit management (where appropriate)

Make sure that your security teams discuss:

- **Standard and planned coding and release cycles** — If the development teams plan on doing CI, how will the code be centrally stored and managed? Security teams should focus on code scrutiny and auditing the code storage/management platform and tools.
- **Tools in use for development, testing and deployment** — Automated testing suites are ideal, but security teams need to understand the tools the development teams plan to use so that they can become familiar with platform security, logging and privilege/credential management.
- **How security can best integrate with the teams** — Ideally, security teams will have some understanding of development practices, and will know how to write test scripts and infrastructure-as-code templates where applicable.
- **Expected standards and behaviors** — If there are no standards to adhere to, what will the team seek to enforce? Think about standards for secure coding, configuration benchmarks (like CIS and others) and vulnerability scan results (what is acceptable to be released).

In addition, security teams should define policies for components, networks and architecture where they can. In other words, they should ask: Where can security create policies that are embedded and applied automatically? Examples might include:

- Configurations for instances and images used in development and production
- App deployment and automation security
- Expected and accepted standards (What does a successful and secure component or deployment look like? Start with the end in mind to ensure you have a target goal.)

One additional area of IT that will likely need to adapt is change management. In traditional IT environments, change requests are often created for weekly or biweekly change windows, where IT staff make changes during the scheduled times (usually off-hours). In a fast-moving cloud application environment, much more rapid changes will need to be allowed. Teams will usually need to adapt by deciding ahead of time which severity of changes will be allowed to occur without prior approval or review versus those that will need more attention. Collaboration platforms can also be useful for enabling more rapid discussions about proposed changes as needed.

“A sound governance model that includes collaborative discussions about code quality, system builds, architecture and network controls, identity and access management, and data security is critically important to developing the standards for controls and security posture.”

Security in the CI/CD World

When integrating into a cloud-focused application development model, security teams need to focus on the following:

- **Code security** — How is code being scanned for vulnerabilities?
- **Code repositories** — How is code being checked in and checked out, and by whom?
- **Code repositories** — How is code being checked in and checked out, and by whom?
- **Automation tools** — What tools are in use to automate builds, deployments, etc.? How can security integrate with these?
- **Orchestration platforms** — How are orchestration tools being used to coordinate and automate infrastructure and cloud components?
- **Gateways and network connectivity** — How can the teams ensure secure connectivity to the cloud for deployments?

Authentication/authorization and privileged user monitoring and management are critical, too. While this sounds obvious, cloud application development pipelines tend to include high-privilege users doing lots of activities, and overallocation of privileges can quickly become an issue without oversight and planning.

When planning for cloud application development, security teams first need to work with application development groups to perform threat modeling and risk assessment for the deployment types that they envision. By performing a threat modeling exercise, security and development teams can better understand the types and sensitivity levels of the assets they protect, how to manage and monitor them in the cloud, and the most likely threat vectors for those assets. The type of data that is stored, transmitted and processed makes a difference when assessing the risk of systems and applications in the cloud. Some data types dictate specific security controls, as well as provisioning into compliant cloud provider environments. Risk assessment and analysis practices should be updated to continually review the following:

- Cloud provider security controls, capabilities and compliance status
- Internal development and orchestration tools and platforms
- Operations management and monitoring tools
- Security tools and controls, both on premises and in the cloud

After risk reviews, and keeping the shared responsibility model in mind (meaning cloud providers and consumers share responsibility for security at different layers of the stack), security teams should have a better understanding of what controls they currently have, what controls they need to modify to successfully operate in the cloud, and what the most pressing concerns are (as they change). It's almost a guarantee that some security controls—tools, processes, policies, etc.—won't operate the way they did on premises, or won't be available in cloud service provider environments in the same format or with the same capabilities.

Security for the CI/CD Pipeline

In the modern CI/CD pipeline for cloud application development and deployment, one of the most pressing needs for all teams is automation, far beyond what we've traditionally seen in enterprise data centers. With cloud deployment moving faster than ever, security and development teams need to automate static code security scans, dynamic platform build and QA application and vulnerability tests. They also need to automate most (if not all) configuration and operations tasks, including web application firewall (WAF) deployments and network access controls (NACs).

“When planning for cloud application development, security teams first need to work with application development groups to perform threat modeling and risk assessment for the deployment types that they envision.”

For cloud deployments, all application development teams, as well as security teams, also need to embrace API integration/use. Providers like Amazon Web Services (AWS)¹ operate a completely software-based infrastructure that may offer sophisticated APIs for creating workloads, adding security controls around those workloads, updating and integrating new code and images for containers, and much more. In keeping with the theme of automation, scripted and programmatic methods of automating deployments need to make heavy use of provider APIs.

Security teams have a number of security controls and areas of emphasis to consider for all phases of the application development and deployment pipelines, as shown in Figure 1 and discussed in the following sections.

Code/Develop

Ideally, your organization already follows secure coding practices. Security and development teams need to discuss standards for languages and frameworks to make sure risk is acceptable before deployment. This objective can be a tall order, and secure coding and development practices are still not all that commonplace today. Look into static code analysis tools, and ensure the code is secured within repositories:

- Are check-in and check-out procedures defined?
- Do solid role-based access controls exist?

¹This paper mentions the names of products and services to provide real-life examples of how security tools can be used. The use of these examples is not an endorsement of any product or service.




Figure 1. Phases of Application Development and Deployment Pipelines

Cloud providers often have options available for code storage and management that include authentication with strong identity management and robust logging/tracking of activity. AWS CodeCommit is a fully managed source control service that hosts secure Git-based repositories that encrypts all files both in transit and at rest, integrates with AWS Identity and Access Management (IAM) for controlling privileges and access to code stores, and logs all activity in AWS CloudTrail. Additionally, AWS CodeCommit has a wide range of APIs available that can enable automation and integration with third-party static code analysis tools for code analysis and review by security teams. Code can be automatically scanned upon check-in, and bug/vulnerability reports can be sent automatically to the appropriate teams.

Build

Building code and workload stacks for cloud applications should incorporate automated and intelligent security controls as well. This stage should include:

- Validated code
- An approved build architecture and controls
- Automated build testing for compiled code

Above and beyond the aforementioned automation and security controls and processes, we need automated reporting that goes to the proper parties for review. This is what will ultimately contribute to a more effective vulnerability management program across the environment. Much like the previous phase of development (code/develop), the build phase can often be securely implemented within cloud provider environments.

AWS CodeBuild is a fully managed CI service that compiles source code, runs tests and produces software packages that are ready to deploy. Managing encryption of build artifacts is critical, and AWS CodeBuild integrates with AWS Key Management Service (KMS). AWS CodeBuild also integrates with AWS IAM for control over privileges to builds and compiled code, and all activity is also logged to AWS CloudTrail

Package

Packaging is the phase of application development when the build is updated with additional software packages, some of which may be open source or from in-house repositories. It is important for development and security teams to audit open source modules for flaws, then discuss methods to protect code repositories automatically. A regular schedule for threat and vulnerability updates with the development and operations teams should be decided upon and incorporated into defined processes.

“Security and development teams need to discuss standards for languages and frameworks to make sure risk is acceptable before deployment.”

Some traditional vulnerability scanning vendors have adapted their products to work within cloud provider environments, often relying on APIs to avoid manual requests to perform more intrusive scans on a scheduled or ad hoc basis. Another option is to rely on host-based agents that can scan their respective virtual machines continually or as needed. Ideally, systems will be scanned on a continuous basis, with reporting of any vulnerabilities noted in real or near real time. AWS Systems Manager can be used to manage package repositories and secure build images with up-to-date patches and libraries. Tools like Trend Micro Deep Security can help to automate application protection and package validation for workloads, too.

Test

The testing phase is one that can be highly automated. Consider both static and dynamic tools, depending on builds. Keys for security teams during the testing phase are:

- Run security testing that's as seamless as possible (avoid interfering with QA if you can help it).
- Define test cases and tools.
- Define acceptable outcomes that meet policy.
- Automate tools and teach developers/QA engineers to run them.

The last point is a crucial one—security teams need to hand off tools to the application developers wherever possible and not insert themselves into every process.

Involvement is key, but running test tools is something the application teams can do. Security should only perform pen tests and continuous monitoring activities regularly once policies and standards are defined.

Using open source build testing tools like Test Kitchen and Vagrant can simplify internal policy validation before you push them, and also in an ongoing fashion.

To coordinate penetration tests and routine checks to validate policies' effectiveness, ask:

- Are only required ports open?
- Are credentials secured?

- Are encryption keys secured?
- Are privileges assigned properly?

Really, any specific elements of your configuration standard or expected posture should be continually validated and assessed using automated orchestration tools and platforms. Many third-party dynamic application scanning and pen testing service providers have fully integrated into the cloud. These tests can be run upon build check-in, image update or manually as needed, with fully automated reporting sent to the right teams.

Deploy/Upgrade

In this phase, security teams are focused on:

- **Documentation** — Note any bugs that are outstanding; document plans to fix and when.
- **Communication** — Coordinate with development and operations teams to instantiate any controls needed for remediation or stopgaps.
- **Life cycle** — Ensure an approved policy for bug remediation is in place and monitored for future release cycles.

Even though you'll still have bugs, make sure to fix any of a certain severity before you push applications and systems out the door.

Deployment involves more on the operations side. Ideally, controlled and automated deployments will be coordinated and controlled by operations with input from the application development teams involved.

Where does security fit?

- Nothing new is added/changed once approved builds are ready.
- Deployment is done to the appropriate location/endpoints.
- Deployment is performed over a secure channel for cloud (TLS/SSH).
- Checks exist to ensure a failed deployment rolls back.

It is critical for security teams to be invested and involved in the development stage. Secure network channels should be established for any deployment activities, which likely involves the use of dedicated circuits like AWS Direct Connect, VPN tunnels using IPSec and/or secure certificate-based HTTPS with strong cryptographic TLS implementations. Image validation—which will heavily rely on automation and a combination of vulnerability scanning and host-based agents that can validate all libraries, binaries and configuration elements used in the application workloads—should also take place at this phase. Orchestration engines are useful for some of these tasks, as are cloud-native tools like AWS OpsWorks that can reliably and securely handle the configuration and assessment of application images.

Operate

This final stage primarily focuses on protection of applications with tools like NACs and WAFs, as well as monitoring, logging and alerting. Define security use cases for production operations by answering the following questions:

- What events should trigger alerts?
- What events should trigger automated remediation?
- What event severities should be in place?
- What controls are needed to properly secure the environment?

For starters, teams should define deployment attributes that can be monitored continuously.

Examples of quick wins for monitoring include the following:

- Types of instances allowed to be deployed (size and build)
- Image expected for deployment
- Location/source of deployment (such as IP address or account/subscription)
- IAM or other user invoked in operations

These attributes should all be known and relatively inflexible, and can easily be used as simple trigger points for alerting or even automated rollback or preventative actions. For example, if an instance type of m1.small is deployed, and the only approved type is t2.micro, this trigger could cause the workload to shut down entirely. Cloud-native or third-party web application firewalls like AWS WAF can easily be set up to block malicious application attacks like SQL injection, cross-site scripting (XSS) and others.

In addition, they can perform manual or automated blocking of IP addresses based on threat intelligence that incorporates reputation analysis. WAFs can generate detailed logs, too, which security teams can then stream back to a central analysis engine like a SIEM platform.

Best Practices

To summarize, Table 1 describes the key security areas of focus in the modern cloud application development pipeline.

Table 1. Considerations for Key Security Phases		
Phase		Focus
 Code/Develop		Look for static code analysis tools that are in place and performing (ideally) automated code scans for checked-in code. Reports from these scans should be sent to stakeholders that include security teams and/or application developers.
 Build		Tools like Jenkins can be used to create builds, and they often have many plug-ins and local controls that should be tuned. What types of builds are allowed, and where are the images stored? A secure location where image security and integrity are controlled is paramount for this phase.
 Package		Code will need to be packaged for installation on builds, and this should be done through automated tools that also have the appropriate permissions and access controls (keys to check out code, for example).
 Test		The test phase should include Dynamic Application Security Testing (DAST) tools, as well as (possibly) traditional network vulnerability scans and various flavors of pen tests.
 Deploy/Upgrade		Only approved builds with packages/software that passes testing should be deployed over a secure channel.
 Operate		Now we're in operations, where we should have "guardrails" set up like the appropriate account/subscription separation, IAM policies, network controls and logging/monitoring.

Additional Development Security Concepts for Cloud

Along with core security controls and practices in each major phase of a modern development pipeline, some additional topics and concepts should be in place. Think of these as overarching concepts that apply throughout the entire life cycle. Figure 2 illustrates these concepts, which we cover in the following sections.




Figure 2. Additional Security Considerations Throughout the Life Cycle

Secrets Management

A critical aspect of managing security in a cloud environment is to carefully limit and control the accounts and privileges assigned to resources. All users, groups, roles and privileges should be carefully discussed and designated to resources on a need-to-know basis. The best practice of assigning the least privilege model of access should also be applied whenever possible. Any privileged accounts (such as root and the local administrator accounts) should be monitored closely—if not disabled completely or used only in break-glass procedures.

In addition to privilege management in configuration definitions, application development teams need to ensure no sensitive material like encryption keys or credentials is stored in definition files, on systems that are exposed or in code that could be exposed. As encryption and data protection strategies are increasingly automated along with other development activities, it's critical to make sure the proverbial keys to the kingdom are protected at all times. In the cloud, this can be easily accomplished with a variety of tools like AWS Key Management Service (KMS) and AWS Secrets Manager.

API Security

As mentioned earlier, APIs are integral to building a robust and automated development pipeline. The security posture of APIs should be documented by providers, and all APIs should be strongly controlled through IAM policies. Use of APIs should be carefully monitored, too, with full logging to AWS CloudTrail and other logging engines.

Privilege Management and IAM

Strong privilege management is a necessity in fast-moving application pipelines. Integration with secrets management tools and a granular IAM policy engine like AWS IAM is crucial, along with federation capabilities and integration with directory services. Security teams should help to define the appropriate least privilege access models needed for all stages of application development and deployment, and then implement this in a centralized tool/service whenever possible. A fragmented privilege management and IAM implementation strategy often leads to poor operational oversight of users, groups and permissions, so a single policy engine should be used if at all possible.

In addition to these overarching technology concepts, some newer technologies are also being heavily used in application development and deployments today, including containers and serverless applications, discussed next.

“Application development teams need to ensure no sensitive material like encryption keys or credentials are stored in definition files, on systems that are exposed or in code that could be exposed.”

Containers and Container Management/Orchestration

Containers are rapidly becoming a common means of quickly deploying application workloads in both internal and cloud environments. Containers are created on a shared OS workload, and both the runtime container image and the underlying OS platform need to be secured and maintained much like other images described earlier. Having a secure repository for container images like Amazon Elastic Container Registry (ECR), as well as orchestration tools that can be used for starting, stopping and managing container deployments securely like Amazon Elastic Container Service (ECS) and Amazon Elastic Kubernetes Service (EKS), is important for enterprises using containers in the cloud. Encryption and IAM controls for images, as well as strong logging for all activities should be priorities.

Serverless Applications and Security

A final type of technology that many application development teams are employing is serverless, which offloads the entire workload (container and OS instance) to the provider's backplane, allowing developers to create microservices applications that only require application code to be uploaded and operated within the cloud provider environment. Serverless security should involve static code review (numerous third-party providers can integrate into serverless environments like AWS Lambda to scan the code), privilege and permission control over all serverless applications with IAM, and complete logging of all serverless application updates and execution using tools like AWS CloudTrail.

Use Case

For modern hybrid application development pipelines, security needs to be integrated in a number of places. Imagine a fictional organization, ACME Corporation, that needs to integrate security into its hybrid cloud application pipelines with both on-premises resources and those running in AWS. Internal code repositories are synchronized from on-premises code repository tools with AWS CodeCommit across an AWS Direct Connect channel, where all code is encrypted and protected with strong IAM policies that restrict code access and updates to a limited team of developers. All code updates, check-ins and check-outs are logged and recorded in AWS CloudTrail. A third-party static code analysis tool is integrated into AWS and automatically scans all code that is updated and checked in. Reports are automatically sent to security and development team members to review the criticality of bugs discovered for remediation.

AWS CloudFormation templates are used to create builds with approved Amazon Machine Images (AMIs) and container images stored in the Amazon ECR, which is also carefully controlled through IAM policies. In the build and update phases, a dynamic vulnerability scanning platform with agents and network scanning capabilities is integrated to scan all application builds for libraries, binaries

and OS configurations to ensure no vulnerabilities are present before deployment. Reports are again automatically generated and sent to team members for review. If the reports show that all images meet pre-approved standards, the images are then pushed into deployment with defined orchestration using Amazon EKS and Amazon EC2 instances with AWS Systems Manager installed for monitoring and administration. Once deployed, AWS WAF is enabled to protect applications from malicious application attacks.

Summary

For modern application pipelines, there are a plethora of tools available from cloud providers and third-party companies to help automate strong security controls through the entire development and deployment process. A strong governance structure is critical to ensure all stakeholders are involved and on board with the new tools and processes needed, and security operations teams will need to help define standards for code and images, as well as build strong protective and detective controls in the cloud environment.

About the Author

Dave Shackleford, a SANS analyst, senior instructor, course author, GIAC technical director and member of the board of directors for the SANS Technology Institute, is the founder and principal consultant with Voodoo Security. He has consulted with hundreds of organizations in the areas of security, regulatory compliance, and network architecture and engineering. A VMware vExpert, Dave has extensive experience designing and configuring secure virtualized infrastructures. He previously worked as chief security officer for Configuresoft and CTO for the Center for Internet Security. Dave currently helps lead the Atlanta chapter of the Cloud Security Alliance.

Chapter 9: How to Protect a Modern Web Application in AWS



Shaun McCullough

SANS Instructor

"As organizations transition their workloads into public clouds, a whole host of new attack vectors emerge, along with new tools to protect. Organizations are faced with the daunting task of trying to cover all the vulnerabilities with limited time and budget. In this chapter, I explore how to stand up a threat modeling program that helps organizations prioritize mitigations and understand the changing threat landscape in the cloud by examining real-world use cases."

Introduction

As businesses move more assets to the cloud, having a security plan is essential, but nobody has the time or resources to do everything that is needed from the start.

Instead, organizations need to prioritize their security plans based on the risks to which they are exposed. Too often, organizations start with securing the service they know best or have read about in a blog, or they try to buy their way out of the risks with multiple, expensive security appliances.

While the team is knee-deep in transitioning core services, security takes a back seat. It's confusing to understand where the cloud service provider's responsibility ends and the customer's responsibility begins, or how best to secure the services and leverage new tools properly.

Prioritizing the risks, and hence determining what should be secured first, can be simplified through threat modeling—the process of identifying and prioritizing the risks to infrastructure, applications and the services they provide. A proper threat model allows organizations to identify applicable risks, prioritize those risks and evaluate how to manage changes in risks over time.

Implementing threat modeling in the cloud is similar to implementing for a traditional infrastructure, but the cloud services, risk priority levels and potential solutions can be vastly different. A threat against a web application stack will be the same in the cloud as it is when deployed on premises. However, cloud providers offer new tools to address the risks. Security teams can bring together cloud-native services, centralized logging, new identity access management processes and easy-to-implement third-party services to make applications and infrastructures safer.

This paper is a use case of modeling the threats against a web application server and how to address those risks in a cloud environment. We will cover the web app stack, including the web server, the application code, and the DevOps pipelines to manage it. Database threats will be covered in future papers in this series. We'll examine the tools and services that cloud providers offer to operate web applications at scale and integrate security services. The paper also breaks down the DevOps process, explains how it can be threat-modeled, and describes common security risks and improvements over traditional workflows.

A Threat Modeling Primer

As defined in a special publication by the National Institute of Standards and Technology (NIST), threat modeling is “a form of risk assessment that models aspects of the attack and defense sides of a

particular logical entity.”¹ By implementing a threat modeling process, organizations can improve their security posture, identify unrealized risks and provide their leadership with the proper tools to prioritize which risks to focus on first.

Threat Modeling Process and Frameworks

Most threat models start in one of two ways:

- Identifying a set of attacker techniques the organization is at risk from
- Identifying a set of deployed assets that are at risk

Organizations need to pick the approach that works best for them, but asset-focused threat modeling is usually the most straightforward.

Threat modeling is a process, not a one-time whiteboard session on a Monday afternoon. As the threats evolve, so do an organization’s risk appetite and security implementations, along with the experience of the team. Organizations must create a culture of threat modeling, where the model is evaluated, implemented, tested, reviewed and re-evaluated regularly.

The first threat model an organization builds could take time and even be painful. As the team gains experience, the process becomes more natural and standardized. Security teams should hold quarterly reviews to make updates, question assumptions and adjust risks. Teams should also perform a yearly re-evaluation of the whole threat model, with all the experts available. Regular reviews of the threat model help organizations understand whether the risk-reduction plans are working.

“Building a culture of threat modeling prepares organizations to address the most significant threats with limited resources.”

¹Draft NIST SP 800-154, Guide to Data-Centric System Threat Modeling, <https://csrc.nist.gov/publications/detail/sp/800-154/draft>

Drivers of Threat Prioritization

Prioritizing threats is often tricky and likely influenced by the expertise or culture of the organization. If the network team is seasoned, runs a stable environment and has the time to research new threats, it can create the most detailed plan for reducing security risks in the team's responsibility area. In contrast, a host team caught in the middle of a complicated operating system upgrade has no time to think of next week's risks, much less next year's. The organizational culture, workloads, expertise and maturity drive how organizations respond to threats. A threat model process helps level the playing field by giving the appropriate team members the space, tools and support to think about risks and threats across the organization.

Among the various threat modeling frameworks, the DREAD risk assessment model works well. Used at OpenStack, DREAD helps teams evaluate the potential results of an attack. DREAD helps the team walk through how a system is at risk, what the attack vector looks like, how likely the attack is to occur and how to prioritize which risks to focus on.

The IANS Pragmatic Threat Modeling Toolkit is a spreadsheet that helps organizations walk through the DREAD framework. Users can identify assets at risk, work through DREAD rankings and graph results for easier understanding.²

“Threat modeling is a process, not a one-time whiteboard session on a Monday afternoon.”

²IANS Pragmatic Threat Modeling Toolkit, https://portal.iansresearch.com/media/739278/ians_pragmatic_threat_modeling_toolkit.xlsx

Risk Assessment and Prioritization

Every risk in an environment is addressed in one of four ways, as illustrated in Figure 1.

- **Mitigate**— Putting a firewall in front of your web server will mitigate some attacks, but not all of them. Most security controls focus on mitigating risks.
- **Eliminate**— Eliminating a risk will likely require changing the nature of the asset at risk in such a way that the risk fundamentally goes away. A firewall cannot eliminate all scripting attacks against a web application, but removing all data entry fields and making the website completely static will certainly eliminate whole categories of attacks. Eliminating risks is ideal, but difficult—and usually means re-architecting.
- **Transfer**— When an organization decides to move on-premises infrastructure to a cloud provider, it is effectively transferring asset risks to the service provider. The organization is making a business decision to pay for the provider to manage, secure, provision or operate the service. Cloud providers operate on a shared responsibility model. From a security perspective, that means that parts of the infrastructure stack have been transferred to the cloud provider. It is now responsible for operating, security and managing the assets. Serverless technology is a good example of transferring risk and taking advantage of this shared responsibility model. A customer could spin up virtual machines in the cloud, managing the full stack from operating system to application. The customer is responsible for the patching, configuration and security monitoring of that virtual machine operating system, while the cloud provider is responsible for the virtualization infrastructure, storage and network. Serverless offerings allow the customer to execute a bundle of code, yet have no direct interaction with the executing operating system. The service provider manages the servers in a serverless offering. The risk of operating system vulnerabilities is now transferred to the cloud provider.
- **Accept**— If an organization is unable to mitigate, eliminate or transfer the risk, then it is accepting that risk. It might be a temporary acceptance to be re-evaluated later. In the threat model process, it is healthy for the organization to understand that accepting risk is a valid option that frees it to plan, prioritize, and dive into the other risks.




Figure 1. Risk Management Strategies

As an organization gets more comfortable with its threat model process, it should start incorporating the model into the beginning of the development cycle, helping to identify risks that need to be mitigated or eliminated before the organization has invested the time in creating and deploying it. Include the whole team when modeling a set of services. The developers likely can suggest and implement ways to significantly reduce the risk scores.

Building threat models for IT-operated application services will help with prioritizing and accepting risks. Cloud services offer new opportunities for customers to mitigate, eliminate or transfer those risks for traditional IT service applications and to establish new workflows for developing and deploying those systems through DevOps.

DevOps with Security

DevOps is a process that enables close coordination between development and operation teams.³ That integration enables organizations to develop and quickly deploy new services with zero downtime and improved reliability. The process is especially beneficial for organizations that deploy new versions of software multiple times a day.

“Building threat models for IT-operated application services will help with prioritizing and accepting risks.”

³NIST SP800-190, <https://nvlpubs.nist.gov/nistpubs/specialpublications/nist.sp.800-190.pdf>

To incorporate DevOps, organizations rework testing and deployment processes to be safe, automated and executable at any time. Continuous Integration is the process by which software changes from multiple developers are integrated into a single stack, likely multiple times a day. With Continuous Integration, security teams can avoid the big end-of-a-sprint integration sessions that cause delays and waste resources. Continuous Deployment is the process of building software to be releasable into production at any time, with an easy push of the button. Continuous Integration and Continuous Deployment (CI/CD) require organizations to rethink their planning, development and deployment pipelines to be highly automated. See Figure 2. With CI/CD, every evaluation, decision, configuration or security test that can be automated is automated. If these processes cannot be automated, then the development team must rework the architecture.

DevSecOps takes the DevOps process and builds in automated security evaluation gates. The “Sec” of DevSecOps requires the organization to establish security policies for the product before development starts, implementing them in the testing and deployment pipelines. Automated tests are security policies that become reality, not just words in a binder. The best CI/CD processes incorporating DevSecOps give developers the tools to test the security of their code at their workstations—at the beginning of the process rather than waiting until the end of development and being surprised.⁴

“Companies using on-premises environments have been leveraging DevOps processes to create close coordination between the developers, who create new applications, and operations, which provides the virtual machines they run on. The cloud brings a whole host of services to automate all aspects of the infrastructure deployment and management that on-premises services are unable to match.”

CI/CD is usually focused on deploying applications automatically and continuously. However, the cloud opens a whole new area, allowing the automatic provisioning and deployment of core infrastructure itself. The cloud provides APIs, development kits and specialized services that let customers control every aspect of the infrastructure with DevOps-like processes and tooling.

⁴ Accelerate: Building and Scaling High Performing Technology Organizations, by Nicole Forsgren, Jez Humble and Gene Kim (IT Revolution, 2018)




Figure 2. Continuous Integration and Continuous Deployment (CI/CD)

Imagine creating an infrastructure pipeline where a configuration file is used to build a web application stack. And say that a new version of the web server is released with a software patch, and you want to deploy it. After testing it locally, the team updates the configuration file and checks it into version control, and a CI/CD pipeline kicks in and replaces all deployed web servers with the updated versions—automatically.

CI/CD comes with risks, however. Automating processes traditionally done by humans can reduce errors, but it also hides unforeseen problems. The platforms that implement DevSecOps and CI/CD pipelines are new attack vectors. The CI/CD platform must become part of the threat modeling process for an organization to ensure that the entire infrastructure is evaluated.

Threat Modeling a Web Application

As previously discussed, the threat model process starts with identifying deployed assets that are at risk—assets that are well understood and vital to the business. As part of our use case, let's model the threat to the web application itself and investigate a threat model for the web application.

Risk of Web Application Attacks

Web applications are usually at risk—they live on the internet, with the sole purpose of capturing and providing information to all their users living on untrusted networks. Complex web applications with user access controls, database-backed pages and free-form input fields are notorious for their vulnerabilities.

The Open Web Application Security Project (OWASP) Top 10⁵ is the best starting place when analyzing

⁵ OWASP Top Ten Project, www.owasp.org/index.php/Category:OWASP_Top_Ten_Project

threats against web applications. Top attack techniques are prioritized, researched and documented, with details of how the attack works and suggested best practices for stopping the attacks.

Cross-site scripting (XSS) is a common attack on web applications that the OWASP Top 10 – 2017 report describes:

- XSS flaws occur whenever an application includes untrusted data in a new web page without proper validation or escaping, or updates an existing web page with user-supplied data using a browser API that can create HTML or JavaScript. XSS allows attackers to execute scripts in the victim's browser which can hijack user sessions, deface web sites, or redirect the user to malicious sites.⁶

Use Case: Spoofing an Identity

Web applications require data inputs and dynamically display information back to users. XSS could result in many different threat categories. For this use case, an XSS attack that exposes other users' browser session credentials can be used to spoof an identity.

After categorizing the threat, a team can evaluate the risk using the DREAD model. Each DREAD risk-rating category is given a value from 1 to 10. Figure 3 describes the ratings.

The rating of a single threat does not provide a full picture of the organization's vulnerable landscape. DREAD ratings of multiple risks should be viewed in tandem to get a complete picture of the risks that need to be prioritized. While informed by the DREAD rating guidance, organizations will arrive at their final rating number/prioritization through a combination of the ratings and their own experiences, knowledge and biases. Table 1 shows the DREAD rating for our use case.

Because XSS is a well-known and well-researched attack method, security teams have multiple ways to mitigate the risk of an XSS attack on a web server. A popular security control is incorporating a web application firewall (WAF) to monitor and block any suspicious traffic before it reaches the web server.⁷ Large cloud service providers make it easy to implement a WAF right from the console. AWS's WAF service allows you to customize rules and access control lists to fit your business and risk models.

⁶The 10 Most Critical Web Application Security Risks, [www.owasp.org/images/7/72/OWASP_Top_10-2017_\(en\).pdf](https://www.owasp.org/images/7/72/OWASP_Top_10-2017_(en).pdf)

⁷Web Application Firewall, www.owasp.org/index.php/Web_Application_Firewall

Damage Potential—How much damage will occur if this vulnerability is compromised?

- 0 = None
- 3 = Individual user data is compromised or affected, or availability is denied
- 5 = All individual tenant data is compromised or affected, or availability is denied
- 7 = All tenant data is compromised or affected, or availability is denied
- 7 = Denied availability of a component/service
- 8 = Denied availability of all components/services
- 9 = Compromised underlying management and infrastructure data
- 10 = Complete system or data destruction, failure or compromise

Reproducibility—How reliably can the vulnerability be exploited?

- 0 = Very hard or impossible, even for administrators; the vulnerability is unstable and statistically unlikely to be reliably exploited
- 5 = One or two steps required; tooling/scripting readily available
- 10 = Unauthenticated users can trivially and reliably exploit using only a web browser

Exploitability—How difficult is the vulnerability to exploit?

- 0 = N/A We assert that every vulnerability is exploitable, given time and effort; all scores should be 1-10
- 1 = Even with direct knowledge of the vulnerability, we do not see a viable path for exploitation
- 2 = Advanced techniques required, custom tooling; only exploitable by authenticated users
- 5 = Exploit is available/understood, usable with only moderate skill by authenticated users
- 7 = Exploit is available/understood, usable by non-authenticated users
- 10 = Trivial—just a web browser

Affected Users—How many users will be affected?

- 0 = None
- 5 = Specific to a given project
- 10 = All users

Discoverability— How easy is it to discover the threat, to learn of the vulnerability? (By convention this is set to 10 even for privately reported vulnerabilities).

- 0 = Very hard to impossible to detect even given access to source code and privileged access to running systems
- 5 = Can figure it out by guessing or by monitoring network traces
- 9 = Details of faults like this are already in the public domain and can be easily discovered using a search engine
- 10 = The information is visible in the web browser address bar or in a form

Figure 3. DREAD Risk Ratings⁸

⁸ Adapted from DREAD Rating, <https://wiki.openstack.org/wiki/Security/OSSA-Metrics#DREAD>

Table 1. DREAD Rating for Web Application

Category	Rating	Spoofing Identity
Damage Potential	2	The business unit is a significant driver of the risk rating for an application. What data does the application hold? How far-reaching would the attack be? How important is the asset itself? In this example, an XSS attack to gain credentials does not do any damage itself.
Reproducibility	7	Once identified, an XSS attack is easy to reproduce through scripts. Only common application access is necessary, rather than special access privileges.
Exploitability	4	Depending on the vulnerability of the application, an XSS could be easy or hard to exploit. Discoverability rates how easy it is to determine if there is potential for an XSS; however, making the exploit perform the desired identity spoofing can be tricky, so we will rate this lower.
Affected Users	4	An XSS attack affects the users logged into the application at the time of the attack, and potentially any users who view the corrupted data. Some users will be affected, but not all.
Discoverability	7	Entering JavaScript into a webpage and reviewing the results gives an attacker a good idea if there is an XSS vulnerability, even if they cannot complete the exploit.
DREAD Average	4.8	

Larger cloud service providers may offer WAF assets that can be integrated into their service offerings. They are easy to set up, are relatively inexpensive, and should be able to block OWASP Top 10 and other common attacks. If the DREAD risk is higher and more protection is needed, the cloud service provider often has a variety of top-tier third-party products with WAF offerings available for installation (for example, Imperva SecureSphere and Fortinet FortiGate).⁹ One way to eliminate the risk of XSS is to remove data entry fields altogether. It requires rethinking the web application architecture and possibly removing functionality for the sake of security. If eliminating the data entry fields is not viable, you can transfer that ownership to a third party. For instance, if the data input fields are for user authentication, leverage a third-party single sign-on service. Eliminating and transferring risks tends to be more costly, but will help decrease DREAD risk scores. The bottom line is that the threat modeling process should drive prioritization of assets and financial commitments.

Use Case: SQL Injection Attack

Modern web applications are driven by databases that can contain a wealth of knowledge that attackers want. A SQL injection tricks the database into returning unintended data.¹⁰ One outcome of a SQL injection attack is information disclosure.

⁹This paper mentions product names to provide real-life examples of how varying classes of tools can be used. The use of these examples is not an endorsement of any product.

¹⁰SQL Injection: Modes of Attack, Defence, and Why It Matters, www.sans.org/reading-room/whitepapers/securecode/sql-injection-modes-attack-defence-matters-23

Table 2. DREAD Rating for Database

Category	Rating	Information Disclosure
Damage Potential	7	A SQL injection, if successful, will likely affect all the data in the database, not just specific users. The actual damage done in information disclosure is another measure that requires the business units to weigh in.
Reproducibility	7	Once a SQL injection attack is identified, it is repeatable.
Exploitability	5	SQL injection (or NoSQL) tends to be easier to accomplish than XSS.
Affected Users	2	Other users may not even notice if a SQL injection attack is happening unless it is damaging the data. For an information disclosure categorized attack, the user effect is nominal.
Discoverability	6	Like XSS, the SQL injection vulnerability is easier to identify than actually to exploit.
DREAD Average	5.4	

The DREAD rating determines the severity of this attack in the environment. See Table 2.

The processes for mitigating a SQL injection and XSS attacks are similar. The SQL injection attack comes through the web application itself; thus the WAF is in a position to identify and block potential SQL injection attacks. Not all SQL injection attacks will be detected, and significant research has gone into countering a WAF.¹¹ When deciding on a WAF product, look at the entire threat model process and ensure that the WAF covers all the threats at the same time.

Another option is to leverage secure coding practices to develop safer code that neutralizes invalid text field inputs before being run in the SQL query on the database. Depending on the programming languages, a number of libraries, design patterns and tools can do this. The security team will need to ensure that all code is following these standards or incorporating the right tools. Today, CI/CD platforms provide opportunities to continuously scan, evaluate or test code as it is being developed.

Now that we've looked at modeling the threat to the web application, let us look at the threat to the development and deployment platform that is used in cloud operations.

Threat Modeling the DevSecOps Platform

We have looked at threat models for a well-known architecture like the web application. Now let's walk through a practical threat model of a CI/CD platform. Again, DREAD helps to prioritize the risks.

¹¹ SQL Injection Bypassing WAF, www.owasp.org/index.php/SQL_Injection_Bypassing_WAF

A CI/CD process is all about safely automating workflows. The Continuous Integration process kicks off when a developer checks code into the designated source code repository.

Distributed version control systems (DVCSs) will mirror an entire copy of the codebase, including all history, on every developer's computer.¹² Git is the most popular DVCS in use today, used with a central Git repository management system like GitHub, GitLab or AWS CodeCommit. When developers request to check their code into the designated central repository, the Continuous Integration system kicks off to test the integration to ensure that it does not break the application. See Figure 4.




Figure 4 Continuous Integration Process

Use Case: Credential Disclosure

Web applications can make database connections directly to query for data. Many times, the web application connects to the database through credentials stored in a configuration file on the application's server. The developers have an instance of the database in their environment for testing, which may include a small copy of production data to test code changes properly.

¹² Getting Started—About Version Control, <https://git-scm.com/book/en/v2/Getting-Started-About-Version-Control>

If that credential file is accidentally checked into the source control system, that configuration file could become visible to unauthorized users—especially with open source software where the DVCS is accessible to the public. Disclosure of credentials can lead to an unauthorized login to the database, called “identity spoofing.” Using the spoofed identity can then lead to additional information disclosure, tampering of data or even denial of service. Identifying each step and categorizing the actions along the way is building up the attack tree.¹³ See Table 3.

As the developer is checking in new code in a Continuous Integration process, it is possible that the developer will accidentally check in that credential file and risk disclosure. If undetected, exposure is guaranteed.¹⁴

In CI/CD, the automated test platform could be used to evaluate the code to look for strings that resemble credentials and reject the merge. These tools are inexpensive and are easy to configure and execute; they fit perfectly with the CI/CD process and will mitigate the credential disclosure risks.

To eliminate the risk of credentials being checked in, eliminate the credential file. Secrets management systems, which are available from cloud service providers or through the marketplace, can be used to programmatically store credentials and only provide them to applications that are authorized. Although this risk-reduction will be harder to implement and can cause changes to the asset, eliminating a risk versus mitigating that risk might be worth the cost.

Table 3. DREAD Rating of Credential Disclosure		
Category	Rating	Credential Disclosure
Damage Potential	5	The damage from information disclosure varies depending on the value of the credentials themselves. In this use case, the credentials at risk are for the development environment and reside on the developer's machine. Because this test database contains a snapshot of production data for testing, customer data is at risk.
Reproducibility	8	The threat exploited is highly reproducible because the attacker can log into the at-risk asset.
Exploitability	8	Logging in with unauthorized credentials is easy when you have the credentials.
Affected Users	5	The database at risk in this particular threat model is a developer's test environment with limited production data.
Discoverability	9	The software is continuously scanning source code repositories looking for credential-like data, thus discovering the data could take mere minutes.
DREAD Average	7	

¹³ Attack Trees, www.schneier.com/academic/archives/1999/12/attack_trees.html
¹⁴ I accidentally pushed sensitive info, <https://github.community/t5/How-to-use-Git-and-GitHub/I-accidentally-pushed-sensitive-info/td-p/225>

Use Case: Software Vulnerability to Denial of Service

Humans write software, and humans are experts at making mistakes. Security professionals are continually patching, monitoring and managing software updates. To make matters worse, developers are increasingly reliant on software packages distributed by other developers. Code actually written by the development team may be a small percentage of the entire code base for the application. For this threat model, teams must evaluate the risk of a vulnerable third-party NodeJS module making its way into the software stack.

Node Package Manager (NPM) is the most widely used NodeJS package delivery tool, and is likely what organizations are using for JavaScript-based frameworks. A vulnerable NodeJS module can cause information disclosure, escalation of privileges or denial of service.¹⁵

Let's look at denial of service and rate the DREAD risks, as shown in Table 4. It can be difficult to know if a vulnerability exists in any included NodeJS packages. Although the vulnerability may not exist in the packages themselves, each of those packages could rely on other packages, which could be vulnerable. The CI/CD platform must continually analyze deployed modules for vulnerabilities discovered post-deployment.

Table 4. DREAD Rating of Software Vulnerability		
Category	Rating	Denial of Service
Damage Potential	7	The amount of damage caused by a denial of service is a business-unit-led decision. Is this a core part of the organization's business? Could it go down for a day and see no real effects? Business drivers are just as important as security risks in the threat model process. Knowing how vital each service is to the business helps define these values. For this use case, the product is a core part of the business and could not go down for any length of time.
Reproducibility	5	Reproducibility can be difficult because the exploit in the NodeJS module could be easy or hard to implement depending on what it is. Predicting future vulnerabilities is impractical. The threat modeling team will have to decide how to handle these ambiguous ratings and be consistent.
Exploitability	5	Similarly, exploitability is hard to assess.
Affected Users	8	The number of affected users can be significant. Denial of service attacks against production systems may slow down or even stop customers from using the application.
Discoverability	3	Because this use case is not an open source application, it will be difficult for an attacker to discover that an application has a particularly vulnerable NodeJS package.
DREAD Average	5.6	

Some code scanner products are available, usually as scriptable software applications that can be run by any CI/CD platform. Commercial versions provide a wealth of threat intelligence and software analysis

¹⁵ NPM security advisories, www.npmjs.com/advisories

and are able to not only identify reported vulnerabilities but also scan deep into the code itself and identify risky functions or statements. The code scanners should be easy to run with the CI/CD platform. When developers integrate their code, third-party vulnerability scanners could scan before acceptance. After deployment, the entire code base should be tested daily for newly discovered vulnerabilities that can flag to the security team.

Expanding on this idea, the entire deployment system can be scanned before deployment. In a cloud service environment, the configuration of the infrastructure itself can be managed by code, using tools such as AWS CloudFormation or HashiCorp's Terraform. When a configuration is changed, a sample virtual machine can be automatically built, then scanned by vulnerability scanning tools to ensure that no known vulnerabilities exist in the packages. Third-party scanners have cloud-ready services that can be initiated by CI/CD in the cloud. The results can be used by the CI/ CD to determine if a deployment should continue—all automatically.

The risk model can help inform decision makers on whether to use free or commercial solutions. Investigate what additional services and intelligence the commercial products provide, whether they will be easier to implement and operate, and how they might work in the build process. Remember, the risk scores from the threat modeling process and the priorities they uncover can help direct where to focus time and money.

Summary

Start building a threat model process as part of the security culture of your organization and reap the benefits throughout the life of your infrastructure. Focus on identifying the threats, the risks they pose, and the relative business importance to help the organization prioritize where to focus attention and resources. The automation of the integration and deployment processes of applications means security policies need to be identified and implemented at the beginning of the development cycle, not the end.

Threat modeling is a great process for identifying risks. We recommend that any threat modeling process do the following:

- Prioritize risks so organizations know where to focus investment.
- Produce concrete plans to mitigate, eliminate or transfer any risks that will not be accepted.
- Bring security into the beginning of system development rather than at deployment time.

- Create a repeatable, improvable process that is used to make decisions, not just a checkbox.
- Document not just the plan but also the risk-reduction results. A threat model process can help organizations understand how effective they are in planning, monitoring, addressing and measuring risks.

As your threat model process matures, teams can start to evaluate risks in systems before they are even developed. Architectural decisions to eliminate a risk rather than only mitigate it will improve security. and likely reduce overall operating costs. And as automated DevSecOps platforms are brought into the organization's workflow, a whole host of risks can be managed automatically.

Adapt a good threat model process that works for your organization. Constantly re-evaluate, improve and expand the process until the organization can see measured results from planned risk reductions.

About the Author

During his 25+ years of experience, Shaun has spent equal parts in security engineer and operations as well as software development. With extensive experience within the Department of Defense, Shaun was the Technical Director of the Red and Blue operations teams, a researcher of advanced host analytics, and ran a threat intelligence focused open source platform based on MITRE ATT&CK. Previously, he was a consultant with H&A Security Solutions, focusing on analytic development, DevOps support, and security automation tooling. Shaun has authored the brand new SEC541: Cloud Monitoring and Threat Hunting and can be found teaching SEC545: Cloud Security Architecture and Operations on a regular basis.

Chapter 10: How to Protect Enterprise Systems with Cloud-Based Firewalls



Kevin Garvey

SANS Community Instructor

"Firewalls are just as important in cloud deployments as they are in on-premises environments. Cloud-based firewalls build upon the principles of on-premises firewall deployments and enable organizations to streamline the administration and usage of data points.

This chapter explore some of the features of cloud-based firewalls, such as web filtering, IDS/IPS, SSO/authentication support, and deep packet inspection (DPI). Ease of deployment of cloud-based firewalls and exciting advanced features are also reviewed. Your security team and your networking team can learn about efficiencies and visibility gained through cloud-based firewalls. Get the most out of cloud-based firewall deployment by understanding all the features organizations can take advantage of."

Introduction

On-premises perimeter security has been a cornerstone of information security programs since the advent of the firewall. Numerous on-premises guidelines and requirements have been drafted to help information security professionals assess their capabilities against best-of-breed compliance certifications. Now, as more organizations realize the rising demand for, and full potential of, migrating their infrastructure and workloads to the cloud, world-class security is no less essential.

Organizations have been meeting the growing demands for securing on-premises networks and data by utilizing the latest generation of firewalls while employing defense-in-depth solutions throughout the enterprise. As cloud migrations have been ramping up over the last few years, the views on network security devices such as web application firewalls (WAFs) and cloud-based firewalls have evolved as well. Gone are the days of deploying network security devices using on-premises equipment only.

Organizations can now virtually deploy WAFs and firewalls in cloud environments. In many cases, the deployment is as quick as pushing a few buttons, reducing the initial setup time from hours to minutes. Organizational focus can now shift from maintenance of the technology—firmware upgrades, patching requirements and physical replacements—to key security initiatives.

The requirements that apply to securing on-premises networks also apply to securing networks that have migrated to cloud environments—but the cloud provides a fresh approach to the security strategy and changes day-to-day expectations.

In this paper, we review how you can rethink on-premises security capabilities and technologies so that your deployments for cloud environments will be familiar and yet improved. We also look at an example of how an organization can successfully implement cloud-based firewalls.

Cloud-Based Firewalls Provide Familiar Features

Since their inception, firewalls have been critical in securing an organization's perimeter. They are the first line of defense against incoming traffic, and the last line of defense for outbound traffic destined for the internet. For years, stateful firewalls that relied solely on port- or protocol-based filtering were sufficient for most organizations. But because bad actors were able to circumvent this simple firewall setup, firewall admins had to look beyond the blocking techniques of traditional firewalls. As the technology matured, firewall engineers and other security practitioners had the responsibility of implementing firewall rules, investigating firewall security alerts and troubleshooting connectivity issues when normal network traffic was disrupted. The latest generation of on-premises firewalls have highly advanced features, and firewall practitioners will find that these capabilities translate very well to a new generation of firewalls: cloud-based firewalls. Figure 1 shows the evolution of firewalls.

Cloud-based firewalls fill an important role. With the increase in cloud implementations, the perimeter has taken on a different meaning and is not as easily defined. Cloud-based firewalls provide the same type of protection as on-premises firewalls, but they protect cloud-based resources and data. These firewalls allow organizations to extend their security controls to various environments in the cloud, including cloud-to-cloud traffic. They solve the problem of capturing traffic from all ingress and egress points, not only those in on-premises environments, but also cloud-connected traffic. All the new capabilities of cloud-based firewalls, coupled with the transfer of operational responsibility out of the end user's hands, has made cloud-based firewalls part of the forward-looking strategic discussions within IT departments.




Figure 1. Evolution of Firewalls

Firewall Features

While firewalls have developed to include functions that address the ever-changing threat landscape hitting an organization's perimeters, many of these features translate well to cloud-based deployments. In particular, features that allow organizations to gather data and inspect multiple on-premises and cloud perimeters help both security practitioners and operations groups make intelligent decisions. The features shown in Figure 2 and detailed in the following sections are important considerations when deploying cloud-based firewalls.



Figure 2. Features of Cloud-Based Firewalls

Web Filtering

Web filtering allows organizations to mitigate against the risk of user activity that does not align with their acceptable use policies. Many organizations have deployed web filtering to monitor user internet traffic and block websites that they deem a threat to the organization's risk posture. Such blocking can be done organization wide, or a more granular approach can allow specified users or departments to bypass the filtering policies for defined websites. Many users are accustomed to web filtering, particularly if they have mistakenly tried to visit a website that is in violation of company policy.

Cloud web filtering is a new iteration of web filtering that allows organizations to enforce web content policies regardless of users' locations. Organizations can set policies based on whether the user is on or off premises. This type of web filtering affords organizations the flexibility to allow users access to the resources they need to be successful while mitigating against activity outside of the company's risk profile. Cloud-based web filtering can also reduce the need for on-premises web filtering equipment.

“Cloud web filtering affords organizations the flexibility to allow users access to the resources they need to be successful while mitigating against activity outside of the company’s risk profile.”

Network Logging

Traditional firewall configurations can produce network metrics on anything visible to them. Firewalls can give an IT group valuable data points on the activity on the network, from blocked and allowed websites to ports being utilized and the duration of network connections. This data allows network administrators and security practitioners to establish a baseline of what “normal” looks like, so that they can identify when the network is in need of troubleshooting or detect anomalous traffic on the network.

Cloud-based firewalls extend an organization’s monitoring capabilities into the cloud. This lets administrators track cloud-based traffic to and from the on-premises environment, allowing security practitioners to establish a baseline for normal cloud network traffic patterns and to identify incongruous patterns. For example, if a rogue vulnerability scanner were running within the cloud environment, changes from the baseline cloud-based network would be detected, and security practitioners would be alerted so that they could investigate.

IDS/IPS

IDS/IPS is a natural addition to any firewall setup. Both an IDS and an IPS watch for questionable network activity by using signature-based rules that search for predetermined patterns in network activity or by analyzing network traffic to identify deviations from the baseline. An IDS is able to identify anomalous traffic but does not block the traffic, while an IPS blocks traffic based upon a predefined set of rules.

IDS/IPS in the cloud works similarly to an on-premises device. Many IDS/IPS vendors offer cloud-based solutions that security teams can deploy easily to protect against cloud-based traffic. Some vendors allow organizations to connect their cloud IDS/IPS deployment to their on-premises solution so that users have a single, comprehensive view.

“Cloud-based firewalls extend an organization’s monitoring capabilities into the cloud. This lets administrators track cloud-based traffic to and from the on-premises environment, allowing security practitioners to establish a baseline for normal cloud network traffic patterns and to identify incongruous patterns.”

SSO/Authentication Support

Firewalls in the past were siloed from directory stores, forcing firewall admins to administer firewall rules and user roles separately. Cloud-based firewalls have the capability to seamlessly integrate with identity and access management (IAM) technologies such as SSO to make the process of administering user roles as simple as possible.

Because cloud-based firewalls can integrate with existing directory stores, admins have fine-grained control of firewall features through existing SSO technologies. This integration also helps eliminate the security risk of stale login accounts on the firewall. Making sure that IAM policies on a firewall stay fresh as users change roles or leave the organization helps to maintain a strong security posture. Cloud-based firewalls make analyzing and correlating SaaS-based application and other cloud-based architecture network traffic easier by showing admins a more complete picture.

“Many IDS/IPS vendors offer cloud-based solutions that security teams can deploy easily to protect against cloud-based traffic.”

The integration of directory services allows network administrators to transfer the responsibility of reassessing users' access from firewall administrators to the appropriate IAM teams. When deploying cloud-based firewalls, an integration with an organization's directory service offers the same features as an on-premises firewall, eliminating the need to audit IAM concerns in cloud-based firewall deployments.

If an organization has not connected its directory store to AWS, it can utilize AWS Directory Service¹ to reduce the burden of maintaining separate accounts in each firewall cloud deployment.

“Cloud-based firewalls make analyzing and correlating SaaS-based application and other cloud-based architecture network traffic easier by showing admins a more complete picture.”

Deep Packet Inspection

Deep packet inspection (DPI) has been included in firewall deployments for years. DPI investigates network packet headers and data to determine whether a packet contains a malicious payload. If the firewall deems the packet to be malicious, the firewall deals with it by following either built-in rules or custom rules developed by the firewall administrator. The most common use case is to drop or block the packet from proceeding to the next hop. Now that firewalls are commonly built with much more processing power, the worry about DPI introducing significant network latency has fallen away, and DPI has become commonplace.

DPI of cloud traffic is just as important as it is for on-premises traffic. Cloud-based firewalls detect malicious traffic not only as it enters the cloud environment, but also as it traverses the cloud infrastructure. This key component allows AWS users, for example, to use VPC Traffic Mirroring in a multi-account AWS environment, capturing traffic from virtual private clouds (VPCs) spread across many AWS accounts and then routing it to a central VPC for inspection.

¹This paper mentions product names to provide real-life examples of how firewall tools can be used. The use of these examples is not an endorsement of any product.

Ease of Management of Firewalls and Firewall Features in AWS

Many cloud-based firewalls allow network and security teams to expand their current, on-premises firewall capabilities to protect their cloud infrastructure. The beauty of the extension is how seamless it is to integrate these new firewalls into day-to-day operations with little operational upkeep by the admin. The following sections point out some of the key features (see Figure 3) that simplify cloud-based firewall deployments.




Figure 3. Seamless Integration of Cloud-Based Firewalls with Operations

Managing All Firewalls in a Single, Comprehensive View

Firewall administrators in the past had to log into firewalls one by one to deploy changes throughout their perimeters. This process created an enormous amount of administrative work for network administrators and security practitioners. More recently, many firewall vendors have provided a single, comprehensive view, allowing teams to save time by making changes on multiple on-premises firewalls at once. Not only has this change been positive for administrators, but it has allowed teams to analyze traffic patterns from a group of firewalls in one console. It also enables richer search results and faster mean time to resolution for security alerts and network outages. Firewall administrators can take comfort in knowing that they can add many of their cloud-based firewall deployments into existing comprehensive views, allowing for easy data correlation between on-premises and cloud-based network traffic.

Deployment Through AWS CloudFormation

AWS CloudFormation provides a common language for describing and provisioning all the infrastructure resources in your cloud environment. With AWS CloudFormation, you can use a simple text file to model and provision—in an automated and secure manner—all the resources needed for your applications across all regions and accounts. For example, using AWS CloudFormation is helpful for cloud-based WAF deployments and ensures all of them are deployed in a consistent manner, making management of each WAF simpler. With the assistance of a master template, AWS CloudFormation is able to launch WAF solutions for web applications. The default configuration deploys an AWS WAF web access control list (ACL) with eight preconfigured rules, but you can also customize the template based on your specific needs.

Advantages of Using a Third-Party WAF/Firewall in AWS

While AWS offers strong in-house-developed firewalls for each customer to deploy, some customers may find it easier to continue their deployment with their existing vendor ecosystem. This allows the customer to enjoy a comprehensive view of their on-premises and cloud-based firewall, and have a simpler license model with their vendor.

Deploying Firewalls in Hybrid Architectures

Many organizations have operational and security requirements in their on-premises environments that they think cannot be properly met in the cloud. Some of them have decided to pursue an intermediate approach, setting up a private cloud, which co-exists with on-premises and public cloud strategies.

Private clouds require the same oversight as public clouds and on-premises networks. In addition, the network security requirements in private clouds are very similar. Just as in a public cloud, cloud-based firewalls are a necessity in a private cloud, and deployment is similar. But all firewalls—whether on premises, public cloud and private cloud—should report to a single location to streamline log aggregation and correlation.

Managed and Customized Rules

While several "... as-a-service" offerings have hit the market over the last few years, many organizations are finding firewall-as-a-service (FWaaS) to be an attractive option. The reason is that FWaaS takes all the administrative burden—patching and management of the firewall platforms—out of the hands of administrators and establishes a unified policy among all deployed firewalls in an organization.

Vendors offer FWaaS as a solution to merge and unify rules and logs from disparate firewalls while the customer enjoys a “hands-off” experience. It might seem as if deploying firewalls on-premises, in a private cloud and in a public cloud would cause administrative headaches, but in fact, FWaaS can remove unnecessary administrative burdens and requirements. This type of service allows administrators to push through policies for all the firewalls in their purview.

Advanced Features

Like many security technologies, firewalls have matured since their inception, including the introduction of enhanced security in so-called next-generation firewalls (NGFWs). As firewalls continue to develop, newer security features, such as behavioral threat detection and analytics, are being incorporated to make organizations even more secure.

Behavioral Threat Detection

Many cloud-based firewalls have started using more advanced features in recent years and continue to build upon the other features each year. Given the amount of data that modern firewalls collect, it only makes sense to put some of that data into action.

Behavioral firewalls convert those data points already present in firewalls into predictions of deviations from the normalized baseline. Identifying what users are doing outside of their typical tasks is a great first start to detecting insider threats. Cloud-based firewalls extend behavioral threat detection into the cloud, giving insight into what is happening outside of the organization’s on-premises environment. An additional benefit is that insider threats can be contained more swiftly if organizations can link on-premises behaviors to anomalous cloud-based activity.

Next-Generation Analytics

Cloud-based firewalls let organizations see, through aggregated sets of metrics and data points, the effectiveness of their security posture. For example, security administrators can easily find out the number of DDoS attacks their cloud and on-premises firewalls have prevented. Cloud-based firewalls also allow security personnel to see the external traffic hitting their cloud space and the network traffic traversing that cloud space. This visibility helps security teams recognize threats not yet written into an alert.

Support for AWS Services

When deploying cloud-based firewalls in an AWS account, where the logs of the cloud-based firewall and WAF ultimately go is a decision any organization can make. For example, to receive a history of all AWS WAF API calls made on your account, you simply turn on AWS CloudTrail in the AWS Management Console.

Use Case: Deploying a Cloud-Based Firewall

When deploying a whole new cloud infrastructure, integrating cloud-based firewalls within a new VPC will both reinforce the security-first mindset and ensure long-term measurement and growth of the VPC. And of course, having protection against the latest threats hitting cloud environments is critical. Let's examine the approach "Acme Corp.," a fictional company, used to deploy its cloud-based firewall.

After testing the waters of cloud computing by moving nonessential company infrastructure into the cloud over the last few years, Acme started a migration of its critical assets to the cloud. Firewall administrators noticed that they did not have good visibility into the traffic going in and out of some of the VPCs that were being stood up by Acme. More importantly, Acme was blind to the traffic flowing between VPCs. While Acme's on-premises firewalls were deployed with attention to security best practices and were well maintained, cloud-based firewalls were not being provisioned in a similar fashion. Many cloud-based firewalls did not follow the security requirements of the on-premises firewall setup, nor were they reporting to a centralized console for each network, which was an important provision for its security teams. Acme's move to the cloud enabled the organization to realize all of the operational benefits of a cloud-based environment. Acme was excited to accelerate the migration of its existing on-premises assets to the cloud and wanted to make sure the security and administration of its new assets matched the world-class quality it had in its on-premises environment.

Acme wanted to add the logs from all of the provisioned cloud-based firewalls into its log aggregator. While it was technically possible to connect all of the log sources into the log aggregator and create correlations and alerts on the new cloud-based log sources, Acme knew that cloud-based firewalls would facilitate a much easier method of moving forward with the requirement. What Acme found was that by deploying a cloud-based firewall, it could go beyond that, because the cloud-based firewall allowed for a single, comprehensive view into both its on-premises and cloud traffic. That meant it would take less time to investigate firewall alerts from various environments.

Acme also wanted a better understanding of traffic in its cloud. To do that, it needed first to determine the baseline network traffic in the cloud and then to detect anomalies from the baseline and identify

network segmentation requirements. In the cloud, detection of anomalies cannot be port-based, so using some of the newest cloud-based firewall features, such as behavioral analytics and behavioral threat detection, meets the requirements for Acme’s new firewall deployments.

Another goal for Acme was the capability to quickly see whether any anomalous activity in the cloud was connected to alerts in its on-premises architecture. To accomplish that, Acme needed a solution that would put everything under one management console, which would reduce investigation time for both security practitioners and network analysts.

In the end, Acme felt comfortable that deploying the new features in its cloud-based firewalls would satisfy its security requirements. See Table 1, which summarizes the requirements and challenges Acme had to address.

Acme deployed the metered F5 Big-IP Local Traffic Manager (LTM) + Advanced Firewall Manager. Not only did it provide NGFW capabilities such as comprehensive threat protection, granular control and visibility into Acme’s cloud environment, but it also allowed Acme to deploy secure office-to-cloud connectivity and cloud network segmentation.

Table 1. Requirements and Challenges	
Requirements of Cloud-Based Firewalls	Challenges
Behavioral analytics	Not seeing all traffic moving from on premises to cloud Missing cloud-to-cloud traffic
Comprehensive view	Having to log into multiple management consoles to manage firewall alerts
Next-generation analytics	Needing to have top-of-the-line, cloud-based firewall technology options

Summary

Whenever organizations add new network segments, their compatibility with firewalls and other network security equipment is a top concern. Cloud security migrations are the next-generation leap many companies have been looking forward to for years. As a result, organizations need to look at cloud-based firewalls that are able to work in concert with traditional firewalls to secure the organization and the applications and assets it has migrated to the cloud.

Using cloud-based firewalls enables businesses to focus on what makes them great while moving the heavy lifting of infrastructure and hardware support to the cloud. Cloud-based firewalls free up network administrators and security practitioners to focus on their key job requirements by relying on the cloud to take over many of the tasks they had to take on for so many years.

Today's cloud-based firewalls have brought the best of what security practitioners and network administrators love about NGFWs to the cloud, while also expanding the capability to aggregate cloud data points. This data is used smartly in DPI, next-generation data analysis and behavioral analysis. Cloud-based firewalls are no longer just a requirement for network security; they are an integral part of network- and security-based decisions in a cloud deployment.

About the Author

Kevin Garvey is a SANS instructor-in-training for MGT512 and security operations manager at an international bank based in New York City. He has been a cybersecurity aficionado ever since he became interested in computers, but formalized his passion by moving from a career in IT to become a cyber professional in 2013.

Kevin has worked at the New York Power Authority, JP Morgan and Time Warner, contributing and leading efforts to grow new and existing cyber initiatives. He holds a CISSP, GCIH, GLEG, GCFA, GCFE and GSLC.

<https://t.me/learningnets>

Chapter 11: How to Implement a Software-Defined Network Security Fabric in AWS



Dave Shackleford

SANS Senior Instructor & Author

"The software-defined data center (SDDC) has changed the way organizations build and design network objects and components in the cloud, as well as security controls and network monitoring. This chapter introduces you to the variety of network objects available in the cloud, along with security considerations for each. Architecture is important, too, both in designing standalone virtual private clouds (VPCs) and connectivity between on-premises and cloud networks, as well as interconnecting numerous VPCs and large multi-account cloud environments. Enabling cloud-native and third-party monitoring and defensive network controls are covered in this chapter, too."

Organizations are rapidly adopting and embracing the concept of the software-defined data center (SDDC). This is having an impact on many aspects of IT operations, architecture and security. One of the most significant changes is in the design and implementation of hybrid architectures of cloud networking, which necessitates a shift to software-defined network controls, tools and architecture—all of which impact security.

Fortunately, the rapidly maturing public cloud provides a wide assortment of innovative and robust network controls, tools and services that can be readily enabled to create an end-to-end network architecture rivaling those we've relied on in the past. In fact, a number of cloud-native solutions available now are making it much simpler to build and maintain very large and flexible networks that include strong service level agreements (SLAs) from providers, redundancy and high availability, and capable security options.

Third-party solution providers have adapted a number of platforms and services to infrastructure-as-a-service (IaaS) environments, as well, providing enterprises with even more options for building secure network designs in the cloud.

Today, it's increasingly common for organizations to have a hybrid architecture model that requires routing and connectivity between data centers and cloud providers, network access controls (network ACLs or NACLs) at several layers, traffic inspection, and security monitoring capabilities and much more. Many organizations will likely end up using some combination of cloud-native and third-party tools and services as they architect and build network designs for cloud infrastructure. This paper explores these options and includes recommendations on where different controls and strategies may fit best.

Software-Defined Network Security: A Breakdown

As the world of IaaS has matured, all core networking controls were adapted to cloud-native services. For organizations deploying workloads and infrastructure in IaaS environments, cloud-native network controls offer flexible and tightly integrated capabilities that are easy to enable and maintain. Some common network services available natively in cloud environments include:

- **Routers** — Routing in cloud environments may not require an actual routing platform or appliance, but may instead be accomplished through software-defined route definitions and rules that are operationalized within the provider's native fabric.

- **Firewalls (network access controls)** — Cloud-native network ACLs can be used to control and restrict traffic into and out of the cloud infrastructure, as well as between internal workloads and services.
- **Load balancers** — Two of the most significant drivers for deploying infrastructure into public cloud environments are scalability and availability. Cloud-native load balancing tools are highly capable and resilient.
- **Network gateways** — To facilitate connectivity to cloud workloads and services, a variety of network access gateways can be enabled. These can be focused on internet access, private access to on-premises environments or other gateway devices, or access among cloud-defined zones within one or more accounts.
- **Web application firewalls** — Web application development and operation is one of the most prevalent use cases for the cloud, and web application firewalls (WAFs) can greatly aid in protecting these applications against a wide variety of threats. Cloud-native WAFs are tightly integrated into network access paths and services such as load balancing. They also have flexible API-based logging, monitoring, configuration and operations capabilities.
- **Network address translation (NAT)** — For the SDDC, NAT can be performed in a variety of ways. Cloud providers have highly reliable and automated translation capabilities and controls in place with almost no need for management and oversight. Enterprises can use dedicated virtual appliances that afford them more granular control over translation as well.
- **Network monitoring** — For security teams in particular, the ability to monitor network traffic and patterns of communication is critically important in all operating environments. The cloud infrastructure is advancing rapidly to support a wide range of use cases in these areas, with powerful native services that can be enabled quickly and easily.

In addition, third-party network security solution providers have largely adapted their products and services to integrate natively into cloud environments. These changes are often considered to enhance and augment a cloud software-defined network security stack.

Cloud-Native Network and Network Security Controls

To develop and implement a robust network security strategy, a technology stack and architecture should include a defense-in-depth set of controls that help to achieve the following goals:

- Confidentiality of data and network traffic
- Integrity of the network path to ensure no interception or modification of data and workloads is possible
- Availability and redundancy to meet performance requirements

Ensuring strong access controls and application-level attack prevention and detection is also critical in a best-in-class network security design. A strong control stack may look similar to the one shown in Figure 1, starting with the outermost network security controls at the bottom and working inward toward actual workloads and application-tier protection at the top.

In addition to this core network control stack, network security monitoring controls and services must be enabled to ensure a high degree of visibility and introspection into all traffic traversing the cloud infrastructure. The following sections break down each of these control areas in greater detail.




Figure 1. Core Network Control Stack

DDoS Protection

Malicious actors often initiate DDoS attacks in an attempt to flood networks, systems and applications with more traffic, connections or requests than they can handle. Other types of DDoS attacks are more subtle, targeting specific services in ways that cause them to hang or fail. DDoS defense is a must-have control for many organizations.

Amazon Web Services (AWS) offers its AWS Advanced Shield service for DDoS protection.¹ The standard plan is included for all tenants and defends against the most common, frequently occurring network- and transport-layer DDoS attacks that target sites and applications. The advanced plan includes features such as additional capacity for large DDoS events, native integration with WAF controls, forensic and historical reporting, assistance from the AWS DDoS Response Team (DRT), and some cost protection for charges incurred during an attack.

As the outermost layer of a defense-in-depth network protection model, cloud-native DDoS protection services can help to improve the availability and resiliency of the entire cloud network infrastructure immediately.

Virtual Network Gateways for Connectivity

Organizations can incorporate a number of different connectivity models into their network architecture, including:

- **Internet access** — This model provides direct internet access to cloud workloads with no relation to traditional on-premises assets or network infrastructure.
- **VPN connectivity** — Point-to-point connectivity using IPSec between one or more gateways can allow for protected network traffic in transit, often implemented as a site-to-site VPN between on-premises data centers, branch offices and cloud environments.
- **Dedicated circuits** — For organizations that need a hybrid architecture with guaranteed bandwidth and more stability, dedicated circuits are available that establish a point-to-point, always-on network connection between cloud provider environments and data centers.

¹ This paper mentions solution names to provide real-life examples of how cloud security tools can be used. The use of these examples is not an endorsement of any solution.

- **VPC interconnectivity** — Many organizations employ more than one virtual private cloud (VPC) in one or more accounts. A VPC is an isolated network zone that can be divided into subnets. Organizations may choose to peer these VPCs together to create interconnected network zones.

To facilitate these various types of connections, there are numerous software-defined gateways available in cloud environments. These include:

- **Internet gateways** — Internet gateways (IGs) are basic VPC software-defined objects that allow traffic in and out of a VPC. They can be used to allow connectivity to VPC subnet resources from the internet. These gateways also perform simple NAT operations for VPC workloads. For workload traffic to the internet, the workload's source address is translated to the internet gateway address. For traffic destined for instances from the internet, the gateway translates the address to the destination instance private IP addresses within the subnet. Organizations can set up egress-only gateways for handling outbound IPv6 traffic if needed. It's important to note that internet gateways provide almost no security at all, aside from address translation. They are software objects that are needed to manage traffic operations, but provide little in the way of real access controls or monitoring capabilities.
- **Virtual private gateways** — A virtual private gateway (VPG) is a VPN gateway and a software-defined object that allows IPSec security association (SA) tunnels to be established with another peer. A customer gateway (CG) is the on-premises side of the IPSec tunnel (either a physical or virtual appliance that terminates the other side of the IPSec connection).

An alternative model to a single point-to-point VPN connection is to use a hub tool such as AWS VPN CloudHub, a configuration in which multiple sites can all connect with IPSec to a set of VPGs that use Border Gateway Protocol (BGP) autonomous system numbers (ASNs) in a larger WAN.

- **AWS Direct Connect gateways** — Organizations can establish a dedicated private circuit between on-premises environments and the cloud, or between numerous VPCs (even those in different regions) using AWS Direct Connect gateways. In partnership with numerous WAN and telecommunications backbone carriers, organizations can interconnect numerous physical circuits between both cloud and data center environments with dedicated bandwidth and more flexible quality-of-service (QoS) configuration.

- **Transit gateways** — While organizations can create direct peering relationships for VPCs, numerous interconnected peering arrangements across a larger number of VPCs can be challenging to design and operate. Transit gateways, which are managed through another service called AWS Resource Access Manager (for managing assets across accounts), help teams create a more traditional hub-and-spoke model of network connectivity across VPC peers or AWS Direct Connect circuits.

For performing network address translation, a process known as NATing, a variety of different platforms and methods are available. In addition to the automatic NAT operations that IGs handle, the creation of a dedicated NAT gateway object grants an organization more control over inbound traffic, as well as providing scalability and flexibility in bandwidth and deployment options. For even more control, organizations can use a dedicated instance workload type known as a NAT instance, giving them full control (allowing for numerous inline security controls and services to be enabled on these systems, if desired).

Routing

Because all network elements are software-defined in the cloud, routing definitions can make use of both traditional network definitions (such as IP addresses and subnet designations) and software object references (such as gateway identifiers). Routing is accomplished by creating a set of route definitions, called a route table, that are implemented directly within the cloud provider fabric. In many ways, routing within the cloud is vastly simpler than using traditional complex LAN and WAN routing models and protocols.

“Routing is accomplished by creating a set of route definitions, called a route table, that are implemented directly within the cloud provider fabric.”

Load Balancing

Load balancing is a critical element of all network designs. Load balancers aid in ensuring that availability and resiliency goals are met for all network and application traffic throughout the global cloud

ecosystem. AWS Elastic Load Balancing distributes incoming app traffic across multiple Amazon EC2 instances. Cloud-native load balancers are more capable and flexible than ever. Cloud providers' native load balancers can route traffic based on simple application or network information, and these network-oriented approaches are best suited for standard network traffic or cloud environments that have more traditional application deployments.

These are also good for internal load balancers on the back end, distributing traffic to storage nodes. Platforms like AWS Elastic Load Balancing can also route traffic based on advanced application information that includes the content of the request and more granular microservices architecture. This is the preferred type of load balancer, especially for any internet-facing and web application traffic. AWS Elastic Load Balancing can also establish HTTPS sessions with clients, making the service highly valuable for mobile access and any secure data transmission. Because most web applications move to HTTPS by default, this becomes more and more relevant. You can easily upload your own certificates to cloud load balancers or use certificates from a cloud-native certificate authority (CA).

“Load balancing is a critical element of all network designs. Load balancers aid in ensuring that availability and resiliency goals are met for all network and application traffic throughout the global cloud ecosystem.”

Web Application Firewalls

Many current WAF offerings can help to protect application workloads from common threats, as well as monitor application interaction to highlight suspicious or malicious behaviors. Cloud providers have integrated WAF policies and capabilities into both platforms and services within their fabric, and customers looking to add application-layer prevention and detection capabilities to their cloud network stack can easily enable these policies and capabilities. In addition, organizations can also enable many third-party solutions to provide advanced WAF policy controls.

Virtual Network Security Appliances for Advanced Protection

For enterprise-grade network security capabilities, a third-party service or virtual appliance may make sense for a number of reasons. Mature solution providers can offer more advanced next-generation firewall (NGFW) platforms. For example, these network security gateways can provide access controls, intrusion prevention, malware detection and other functions that enhance and improve a comprehensive cloud networking strategy. Some of these capabilities may also mimic functions available on premises, possibly helping to achieve audit and compliance goals. Leading providers like Fortinet offer numerous cloud marketplace offerings that integrate with IaaS environments to bolster network edge security and allow enterprises to create true perimeter security service zones. These systems and services may afford organizations more deployment flexibility, as well as unified management consoles for hybrid deployment and operations.

Cloud-Native Network ACLs

As a final layer of network defense, cloud-native network ACLs can help prevent attackers from using unapproved network connections to infiltrate systems, moving laterally from a compromised application or system, or performing any illicit network activity regardless of environment.

The first focal area for any cloud-native network isolation and segmentation tool should be the core network zones associated with cloud accounts. In AWS, these are VPCs and can contain any number of distinct network subnets. VPCs can also be peered to one another and connected through AWS Transit Gateways and AWS Direct Connect circuits. Subnets within each VPC can be configured to communicate as needed through routing and cloud-native network ACLs.

Organizations can create and apply cloud-native network ACLs within the VPC to isolate and control traffic flow into the VPC subnets altogether, as well as to and from instance workloads running applications and services. AWS has two built-in types of network access and isolation controls: security groups and network ACLs. Both security groups and network ACLs can control traffic into and out of network deployments. Security groups apply to instance workloads and are stateful, whereas network ACLs apply to VPC subnets and are stateless. Security groups start with a network ACL policy of Deny All, and enterprises can then add rules to allow only those types of network access needed.

Table 1. Security Groups vs. Network ACLs

Security Groups	Network ACLs
Apply to instances	Operate on VPC subnets
Only support Allow rules (layered on a default Deny)	Support both Allow and Deny rules
Are stateful	Are stateless
Are considered in their entirety before traffic is allowed	Are processed in numerical order
Must be associated with an instance to apply	Apply automatically to all instances in a subnet

Cloud Fabric Controls for Network Security Monitoring

In addition to the key network ACLs, another foundation of a sound network security strategy is network monitoring. This has been challenging in the past because the software-defined network fabric of cloud providers didn't have native offerings available to easily monitor network behavior. Further, leading network security vendors didn't have compatible solutions within the cloud for monitoring network traffic, and full packet-capture network "taps" hadn't yet materialized. Fortunately, those issues are now in the past, and network security teams and security operations teams can capably monitor network traffic as needed.

The first type of network monitoring control organizations should enable within the IaaS cloud is collecting network flow data for monitoring communications to, from and between workloads within VPCs. VPC flow logs can be used to monitor and track network events and behaviors at a large scale. With these types of flow logs, customers can designate a storage location for all logs and are also able to aggregate and stream flow logs to SIEM services as needed.

Flow log records include values for the different components of the IP flow, including the source, destination and protocol. VPC flow logs can help security teams in a number of ways, such as troubleshooting and analyzing security group rules, monitoring traffic communicating with workloads, and determining the direction and patterns of traffic to and from cloud network interfaces.

Another capability many network security teams have sought in the cloud is full network packet capture controls. In AWS, a feature called Amazon VPC Traffic Mirroring permits network traffic to be copied from any compatible system in a VPC to a suitable endpoint such as an elastic network interface (ENI), network load balancer and so on. Many network brokering tools and platforms can now leverage this mirroring capability to pull traffic from instances in AWS VPCs, enabling security operations teams to perform deep packet inspection (DPI), network forensics and even selective packet filtering.

Identity and Access Management and Network Isolation/Segmentation

For many organizations, designing software-defined network strategies for the cloud often encompasses a blend of controls that include other services within the cloud fabric. The most common (and important) among these is identity and access management (IAM). Controlling access to network controls, platforms and other network assets within the environment is critical to ensuring that only the appropriate staff and services have access to network configurations. It also improves continuity and stability of the cloud network configuration.

Another key element of IAM that security teams need to adapt to is the use of IAM for isolating assets, enabling teams to create microsegmentation architectures with affinity policies in place. IAM is being used more and more to control access to and interactions with resources in the cloud based on permissions and privilege assignments, making IAM a key factor in access control today. All software-defined assets in the cloud can have policies assigned to them, and this can help manage access just as much as network policy traditionally has, or even more effectively in many cases. Leading cloud providers have a wide variety of prebuilt IAM policies that organizations can enable as service roles, creating strict control models between users and services, users and workloads, services to other services, and really any software-defined object to another within the cloud fabric. Used in conjunction with strong network policies and controls, this can help to improve security and limit the scope of impact that may occur because of a misconfiguration or attack against cloud assets.

In addition to cloud-native controls and services, as well as third-party virtual appliances, we've seen the emergence of a new cloud service model named by Gartner as secure access service edge (SASE), which combines a number of different elements of cloud services and security into a unified fabric.

- **Software-defined WAN (SD-WAN)** — This first element of SASE is oriented toward network access, control and architecture, and allows for interconnectivity between on-premises environments and cloud provider infrastructure through a singular backbone service or vendor solution. These networking services often provide common networking capabilities, such as routing, bandwidth shaping and QoS, and core content delivery network (CDN) services that can set priorities on specific content and service access and transmission.
- **Cloud security-as-a-service (SECaaS)** — This second convergence category included in SASE is a broad category, including services often provided by cloud access security brokers (CASBs) that include data loss prevention (DLP), content filtering, application control,

advanced malware detection and response, cloud provider reputation scoring, user behavioral monitoring and more. In addition, SASE brings together more SECaaS offerings, including VPN replacement technologies, WAF and traditional firewall filtering, network intrusion detection and prevention, and even remote browser isolation (RBI).

In essence, the SASE space looks to take advantage of the cloud brokering model already seen with CASB, CDN and even identity-as-a-service (IDaaS). It includes more networking capabilities and control, as well as combines security services in one brokering model that could potentially simplify the networking and security controls stacks currently in place.

Leveraging Infrastructure as Code for Automation and Guardrails

In the past several years, the concept of using templates to define and manage infrastructure has gained ground. In most virtualized environments, security teams have made heavy use of virtual machine templates and snapshots, and network devices have configurations that they can apply to define a system state. In the cloud, the entire environment is a programmable fabric, providing many more opportunities to implement template-based components and infrastructure objects and services.

This idea, now collectively known as infrastructure as code (IaC), has completely reshaped the way organizations automate and manage infrastructure in the cloud. DevOps teams have embraced this idea for some time, and security teams are beginning to integrate their controls and definitions into these infrastructure templates. IaC tools and templates are available for all major IaaS clouds and can be used to configure and define a wide range of objects and service definitions, including:

- Virtual machines and images
- Container configurations and images
- Network ACLs and subnets/VPCs
- Storage nodes
- Identity policies and roles
- Serverless functions and code

There are many benefits to using IaC, but several are worth calling out explicitly:

- **Reproducible and reusable infrastructure** — One thing that templates can really bring to the table readily is consistency. With a template-based model that defines how security teams want a significant portion of their infrastructures to look, it is possible to maintain all the elements within that template simply by reusing and reproducing the template elements as needed.
- **Version-specific and validated infrastructure** — The entire premise behind IaC is treating infrastructure definitions in templates (and the templates themselves) as we would code checked into repositories. Each template should have a version, and each check-in should ideally have some automation to scan the template for desired configuration elements and included object definitions.
- **Better-documented infrastructure** — There's more opportunity to easily document infrastructure with IaC, because approved personnel can simply add comments and documentation directly into the template files.
- **Infrastructure change monitoring** — Security teams can monitor template files using traditional file-integrity monitoring tools and methods. They now have a much better way to track unexpected or illicit changes to templates that could affect their entire cloud deployment model if not caught.

For software-defined cloud networking, IaC templates can be used to define every control and element of the network stack covered here, including specific access control rules in security groups, route table entries, load balancing configurations, identity policies and much more. Additionally, most third-party tools and platforms can also be referenced and configured through these templates.

Wrapping Up: Best Practices

The software-defined network is simply one critical element of the SDDC. When building a cloud network, consider the following best practices and recommendations for building a hybrid network security architecture in the cloud, using both cloud-native and third-party controls:

- Design an architecture that includes a “transit zone,” where network security access controls and robust intrusion detection can be applied. This zone could be a subnet within a single VPC, a dedicated VPC peered to others or a dedicated VPC that leverages a transit gateway to connect other VPCs and/or on-premises locations. Ideally, a third-party network security virtual appliance should control and inspect all traffic coming into and through this dedicated zone.
- Limit the application of network ACLs to either allow or deny known trusted or malicious IP addresses and subnets. Use security groups to define and apply the majority of the network ACLs. That said, leverage security groups and network ACLs to the full extent needed, because these are capable controls that are wholly integrated and inexpensive to implement.
- If you need more mature and in-depth network security controls (and you likely will), consider a third-party virtual firewall/IPS appliance as a gateway or network security service layer. Build in multiple availability zones and plan for redundancy and failure conditions that may occur. You don't want to be in a position where resources fail and you experience downtime. Use cloud-native or third-party cloud load balancing integrated with APIs and the providers' native fabric within any and all network segments if possible.
- Enable a WAF service, or third-party appliance or service, that provides strong application-tier policies and controls, as well as behavioral monitoring.
- Consider an advanced DDoS protection service plan with your cloud provider if your organization is prone to these types of attacks.
- Enable VPC flow logs and stream them to a dedicated storage node within all cloud accounts. Depending on the volume of flow records generated, a third-party solution for behavioral analysis of flow records may be an additional investment worth pursuing. There are many options available.

- For DPI and network forensics, enable VPC traffic mirroring of important network traffic to a dedicated virtual appliance that can empower the security operations and incident response teams. This powerful capability creates parity with traditional on-premises network traffic capture options like taps.
- Plan IAM roles and permissions to protect access to and use of VPC resources and services. Many VPC objects and services can easily be controlled through IAM, including Amazon Elastic Compute Cloud (EC2), Amazon CloudWatch for monitoring, AWS Elastic Load Balancing for load balancing and much more.
- Wherever possible, make use of IaC templates to define objects and configuration for your network architecture, thus improving consistency and auditability of all controls.

As your cloud environment grows and hybrid cloud network architectures become the prevalent design model, keep these recommendations in mind.

About the Author

Dave Shackelford, a SANS analyst, senior instructor, course author, GIAC technical director and member of the board of directors for the SANS Technology Institute, is the founder and principal consultant with Voodoo Security. He has consulted with hundreds of organizations in the areas of security, regulatory compliance, and network architecture and engineering. A VMware vExpert, Dave has extensive experience designing and configuring secure virtualized infrastructures. He previously worked as chief security officer for Configuresoft and CTO for the Center for Internet Security. Dave currently helps lead the Atlanta chapter of the Cloud Security Alliance.

Chapter 12: How to Build an Endpoint Security Strategy in AWS



Thomas J. Banasik

SANS Analyst

"This chapter provides a high-level overview for designing an endpoint security strategy in AWS. In this chapter, I discuss considerations for traditional versus cloud-based endpoints, integration with SIEM, and response via endpoint detection and response (EDR) platforms. This chapter also explores aligning AWS security solutions to align with existing security investments.

While the target audience is the cloud security architect, these concepts are applicable to cloud security analysts, engineers, and security operations center (SOC) leadership. My intent is to provide a foundation for leveraging endpoint security technologies for secure migrations to cloud-based architectures and zero trust networks."

Introduction

The nature of today's business is driving organizations away from traditional on-premises data centers and into distributed cloud computing environments, and with this move comes the challenge of securing endpoints in a cloud-dominated world.

Not long ago, endpoint security involved little more than signature-based antivirus, but endpoint security capabilities have evolved. Now we have endpoint detection and response (EDR), machine learning (ML), user and entity behavior analytics (UEBA) and data loss prevention (DLP) integrated suites. These cloud-based endpoint security technologies are adapting to industry trends, providing cost-effective, readily deployable and fully integrated solutions to protect assets in the cloud—all managed from a single comprehensive view.

In this paper, we evaluate endpoint security requirements in Amazon Web Services (AWS). We delve into identifying threats, protecting assets, responding to events and recovering from incidents in a distributed cloud environment. This strategy develops a defense-in-depth architecture aligned with organizational business drivers in the cloud. Endpoint security solutions in the cloud provide greater flexibility to manage physical, hybrid and cloud security models while providing enhanced visibility in centralized monitoring services.

Moving Endpoint Security Solutions to the Cloud

The business case for moving to the cloud arises from the economies of scale for computing resources and storage, as physical layers of computing are abstracted to a managed partner. As endpoints are transferred, provisioned or migrated from a physical asset into a cloud model, ensuring their security is critical. A successful endpoint security strategy that addresses the various challenges of cloud migration, such as scale, speed and complexity, can yield better cost savings, visibility, agility and scalability.

Endpoint security solutions in AWS are the hallmark of successful cloud migrations. Amazon Elastic Compute Cloud (EC2) instances provide nearly limitless efficiency gains while encompassing data protection and unparalleled visibility through cloud-native security services including Amazon GuardDuty and AWS Security Hub.¹ AWS also leverages industry-leading partners to streamline tools, ensuring that an organization's defense doesn't blink. These groundbreaking integrations allow security operations teams to identify the indicators of attack (IoAs) and indicators of compromise (IoCs) to act proactively—instead of reactively, after a breach.

¹ This paper mentions product names to provide real-life examples of how visibility tools can be used. The use of these examples is not an endorsement of any product.

“A successful endpoint security strategy that addresses the various challenges of cloud migration, such as scale, speed and complexity, can yield better cost savings, visibility, agility and scalability.”

Importance to the InfoSec Community

Why is an endpoint security solution so critical? With GDPR and its significant penalties for non-compliance, the expectations for data protection have changed. For example, the European Union (EU) holds data controllers and processors responsible not only for personally identifiable information (PII), but also for timely notifications when a breach occurs:

In the case of a personal data breach, the controller shall without undue delay and, where feasible, not later than 72 hours after having become aware of it, notify the personal data breach to the supervisory authority infringements of the following provisions shall, in accordance with paragraph 2, be subject to administrative fines up to 20 000 000 EUR, or in the case of an undertaking, up to 4% of the total worldwide annual turnover of the preceding financial year, whichever is higher.²

Of course, data is stored, processed and accessed via the endpoints that are commonly the user's interface to sensitive data, including PII. Information security starts at the endpoint to build a defense-in-depth architecture capable of securing people, processes and technology. Elevated compliance directives make the endpoint attack vector even more critical in global business operations.

Traditional vs. Cloud-Based Endpoints

What's the difference between traditional and cloud-based endpoints? Endpoints are remote computing devices designed as a human interface to translate data access to and from the network. Traditional

² <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32016R0679&from=EN>

endpoints include laptops, desktops, servers, workstations, mobile devices and the IoT. The cloud environment transfers management of the lower layers of the OSI model—physical, data link and network—to a managed service provider that controls system resources and storage while providing the organization with greater control, agility and security over data.

Defining cloud endpoints is challenging because of hybrid architectures that combine physical, virtual and cloud-based assets. The key to identifying cloud endpoints resides in the service-oriented architecture (SOA) used for providing resources as a service in such models as infrastructure-as-a-service (IaaS), platform-as-a-service (PaaS) and software-as-a-service (SaaS). Cloud-based endpoints include provider-hosted servers, databases, instances, services and applications. Cloud-based endpoint security strategies are designed to secure data at rest, in transit and in use. These technologies include capabilities such as antivirus (AV), a host-based intrusion prevention system (HIPS), application blacklisting, machine learning (ML) and UEBA.

Securing endpoints in hybrid and cloud-based hosting models is very different from doing so in a traditional on-premises data center. With SOA, cloud providers assume shared responsibility for providing resources to customers that are leveraging the cloud's economies of scale. Under that model, the customer is at risk of losing visibility into those cloud resources. Naturally, organizations objected to this, because they require visibility into all of their assets, regardless of where they reside. The traditional data center model leveraged host-based AV and firewalls to secure endpoint data within a defined trust perimeter. The cloud abstracts the concept of on-premises data centers into a decentralized model with a de-perimeterized structure. User endpoints communicate with the cloud network via physical network connections, VPNs, mobile devices and internet-facing web portals. Endpoint communication with management services is critical to enable rapid response for security incidents. While hybrid on-premises security management services integrate with the cloud, best practice recommends leveraging cloud-based SaaS solutions to enhance visibility regardless of where the endpoint lives.

“Best practice recommends leveraging cloud-based SaaS solutions to enhance visibility regardless of where the endpoint lives.”

Use Case: Cloud Endpoint Migration and Integration in AWS

Moving assets to the cloud requires an evaluation of security requirements. This evaluation begins with choosing an endpoint security solutions provider that can provide support in physical, hybrid and cloud-based computing models. After selecting a provider, the organization must review its security requirements to determine which security features, such as ML, HIPS, application blacklisting and UEBA, are required. The organization must establish centralized visibility into assets and then synchronize threat intelligence with the host, as outlined in Figure 1.




Figure 1. Five Steps of Security Endpoint Migration

The due diligence level of this scenario has two key goals:

- 1. Select your endpoint security provider** based on business requirements for protection, migration, time, visibility, consistency, complexity, speed and scalability.

2. Configure endpoint security capabilities to foster integration, and evaluate features including EDR, signature/heuristic-based AV, firewall, HIPS, application blacklisting, DLP, ML and UEBA. Key activities include:

- Evaluating endpoint agent visibility for log sources
- Assessing integration requirements with SIEM
- Testing AV alerting for false positive rates
- Testing HIPS for automation capabilities
- Evaluating UEBA for ease of implementation
- Determining cost savings of ML capabilities

3. Identify assets via cloud-based security managers, and deploy endpoint security agents to physical, virtual and cloud-based assets such as Amazon EC2 instances.

4. Bolster visibility in a comprehensive view service such as Amazon CloudWatch event monitoring, where analysts can easily view endpoint activity.

5. Synchronize threat intelligence with Amazon GuardDuty agentless monitoring and conduct security monitoring in cloud-based SIEM services such as AWS Security Hub.

Endpoint Detection and Response

WEDR agents are a central element of migrating to AWS. Legacy endpoint security products are limited to either blocking or allowing an activity. EDR products add the ability to record endpoint activity and store it for future searches. Capturing IoCs is an ideal feature for integrating EDR agents with threat intelligence services, such as Amazon GuardDuty, which provide continuous threat monitoring and agentless detection for malicious behavior. See Figure 2.

EDR agents also enhance cloud-based security operations by integrating system monitoring capabilities and leveraging system monitor logging and OS equivalents to provide detailed information about processes, connections and file changes. Tracing parent-to-child process relationships is key to determining the root cause of a cyber incident. A traditional security agent might report an endpoint infection, whereas an EDR security agent confirms the threat is blocked and, as shown in Figure 3, identifies the spawning process traced to a recent phishing attack.



Figure 2. Amazon GuardDuty




Figure 3. EDR intercepts the attack cycle before malware spreads.

Signature- vs. Heuristic-Based Antivirus

Endpoint security agents require a robust base of malware file signatures to stop attackers from leveraging known malicious files. Signature detections serve as a baseline of security but are not an assurance of safeguarding data, because an attacker can modify the malware source code in minutes, resulting in a new signature capable of beating signature-based AV. Heuristic- and behavior-based endpoints integrate ML to identify new malware based on behavior instead of signatures.

Application Blacklisting

Endpoint security solutions in the cloud require application control through both whitelisting and blacklisting. AWS Systems Manager and AWS Config provide the capability to record inventory data to enable scenarios such as tracking newly installed or removed software applications, assessing security risk and troubleshooting.³ Endpoint security solutions often include these types of application controls to prevent the use of hacking tools and malicious software. This is often a challenging process because of frequent software updates that change file-based signatures.

User and Entity Behavior Analytics

UEBA is the human equivalent of ML for systems. UEBA leverages the baseline of a user's activity to determine the expected pattern for that user. When a user deviates from the established baseline, or when a user's pattern suddenly aligns with known malicious patterns, UEBA-capable agents trigger alerting and synchronize this data into threat intelligence services such as AWS Security Hub.

Data Loss Prevention

Security teams utilize DLP cybersecurity technology to monitor and alert on data content. This technology supports organizational compliance and data protection requirements for intellectual property, PII and confidential data. DLP technology is a unique solution for PII breach monitoring because of its content inspection capabilities. Cloud-based endpoint security agents with DLP capabilities can alert on the transfer of sensitive data, such as PII or proprietary source code, and alert cyber responders through a centralized monitoring service.

Endpoint Security Solutions in AWS Marketplace

AWS cloud-based endpoint security solutions offer seamless integration. Security solutions currently available in AWS Marketplace offer direct integration with more than 800 security applications from more than 36 leading endpoint vendors. This level of partnership allows organizations to select and integrate the most appropriate endpoint security partner based on business needs and capability requirements. Seamless integration fosters the deployment of endpoint agents across physical, virtual and cloud-based Amazon EC2 instances for total endpoint coverage in the environment.

³ "Preventing blacklisted applications with AWS Systems Manager and AWS Config," April 26, 2018, <https://aws.amazon.com/blogs/mt/preventing-blacklisted-applications-with-aws-systems-manager-and-aws-config>

Amazon GuardDuty allows organizations to take endpoint security further in the cloud through a threat detection service that continuously monitors for malicious activity and unusual behavior to protect AWS accounts and workloads. Amazon CloudWatch provides log visibility to view events and security incidents in greater detail. These capabilities aggregate into a comprehensive view with the AWS Security Hub. Gone are the days of traditional signature-based AV. Today, well-prepared organizations rely on the power of cloud-based endpoint security solutions.

Summary

The flexibility, elasticity and economy of cloud computing are driving organizations to move from traditional to cloud-centric computing models. Cloud migration requires evaluation of business requirements for protection, migration, time, visibility, consistency, complexity, speed and scalability. Cloud-based endpoint security solutions have moved from simple AV to integrated suites capable of securing assets in any environment with advanced capabilities such as application control, ML and UEBA. Synchronization with AWS services such as Amazon CloudWatch for log visibility, Amazon GuardDuty for threat intelligence and AWS Security Hub for synchronization provides a comprehensive view for responders to combat the threat while upholding organizational security objectives in a distributed cloud environment.

About the Author

Thomas Banasik is a SANS analyst and senior security operations center manager for Veritas Technologies, LLC. He has consulted with numerous organizations in cybersecurity across the government, military and commercial sectors. An incident response expert, Thomas has extensive experience in security operations, threat intelligence, insider threat, and threat vulnerability management. He previously worked as a senior security operations center manager for the U.S. Government Accountability Office and is a retired U.S. Army cyber and military intelligence officer. Thomas holds the GCIH, GCWN, GCIA, GSEC, and CISSP-ISSEP, ISSAP, ISSMP certifications and is currently pursuing a second graduate degree in information systems security engineering from the SANS Technology Institute.

<https://t.me/learningnets>

Chapter 13: How to Leverage a CASB for Your AWS Environment



Kyle Dickinson

SANS Instructor & Author

"Cloud service providers offer a range of services, including those where data can be stored. Now how is it that organizations can verify and enforce that their data does not go to an unapproved location or service? Cloud access security brokers (CASBs) can aid in addressing this problem. When security teams want to implement a CASB into the environment, they're able to leverage a phased approach so it doesn't impact end users too much too quickly (unless that's how you want to roll). On top of data loss prevention (DLP), CASBs also monitor user activity through the different integrations to let security teams know if something is "not common" in the environment, so they can gain an understanding if a username/password is compromised."

Introduction

With the explosive rate of enterprises moving toward the use of cloud service providers (CSPs), organizations are seeking new methods and best practices to implement security controls in cloud environments. Many organizations have already securely and successfully migrated their productivity suites and web applications.

Now they are moving their business-critical and highly sensitive systems, including HR applications, customer relationship manager (CRM) systems and enterprise resource planning (ERP) software to the cloud. But how can they ensure that they've gained visibility through and through?

With the convenience provided by cloud access security brokers (CASBs) and the means to integrate with modern technologies, organizations can effectively secure their cloud-based data. In this paper, we discuss ways to integrate CASBs into your organization, common functionalities found within CASB platforms and how CASBs can aid organizations in securing their footprint in the cloud. We begin with CASB deployment types.

Integrating CASBs

CASBs can be integrated into organizations in various ways. It's up to each organization to determine which deployment method best fits its needs. As part of this decision-making process, organizations need to be aware that deployment types differ in the features and functionality they provide. The types of CASB deployments include:

- **API** — This deployment mode, shown in Figure 1, allows organizations to integrate to their applications, and it requires no agents. However, the available APIs are limited to what the SaaS provider allows access to. With the API integration, organizations may not have the ability to do real-time prevention.
- **Forward proxy** — Forward proxies redirect traffic destined to an application to the CASB (see Figure 2). This can leverage agents or proxy auto-config (PAC) files. To leverage a CASB as a forward proxy, organizations must have a solid strategy for managing endpoints because they must elastically deploy either the agent or PAC file. If your organization leverages BYOD, this deployment mode may be challenging to implement.

- **Reverse proxy** — This method redirects traffic through a federated identity to a SaaS application. It integrates with existing identity providers and allows an organization to securely access its cloud applications for managed enterprise devices as well as BYOD. See Figure 3.

Organizations should consider consulting with the CASB vendor to better understand what features are available based on the deployment method they choose to integrate.




Figure 1. API Method




Figure 2. Forward Proxy Method




Figure 3. Reverse Proxy Method

Common CASB Capabilities

Similar to the NIST characteristics of cloud computing (on-demand self-service, broad network access, resource pooling, rapid elasticity and measure service),¹ there's something to be said about the common capabilities available within a CASB solution. Key capabilities that aid organizations in securing their cloud applications and AWS footprint include visibility, compliance, data security and configuration compliance (See Figure 4).



Figure 4.

Comprehensive Visibility

When moving to CSPs and cloud applications, a degree of visibility may be lost, and that may depend on the type of logging available from the provider. Correlating the logs and data available can also prove to be challenging, especially across multiple providers and applications. A CASB assists organizations with the visibility of their cloud-based applications. By giving organizations insight into the security posture of their infrastructure-as-a-service (IaaS) environment, they also gain a better visibility with their SaaS footprint.

Security teams should be looking for patterns of misconfigurations within their Amazon Web Services (AWS) footprint or cloud applications. Examples could include whether there are stale users in the environment or whether cloud storage is allowing anonymous or public access. Teams should also be reviewing whether controls that have been put in place are taking effect and whether there are interactions with unauthorized applications. With an organization adopting multiple services that reside in the cloud, providing a comprehensive view to operations and security teams can reduce the complexity inherent in managing multiple SaaS applications. Because CASBs can integrate with the different cloud applications, including an AWS footprint, they provide security teams with a comprehensive view of the environment.

¹ "The NIST Definition of Cloud Computing," NIST Special Publication 800-145, September 2011, <https://csrc.nist.gov/publications/detail/sp/800-145/final#pubs-documentation>

Auditing

CASBs can help make sense of AWS CloudTrail² auditing data, including how to determine malicious behavior. They typically use machine learning (ML) to baseline normal behavior in an environment to reduce the number of false positives. Although CASBs can also evaluate an organization's AWS footprint, including auditing data, so it has a better understanding of the activity occurring with the AWS footprint. This information is critical because the size of an organization can scale from tens to hundreds of accounts. With a CASB's capability to perform user and entity behavior analytics (UEBA), security personnel can get a better understanding of behavior that deviates from the norm. Consider these examples:

- Is it normal for David from Accounting to stand up multiple SQL databases in a day?
- Should Marc the intern be deleting Amazon EC2 instances after hours?
- Is someone logged in from multiple locations at the same time?

UEBA gives organizations the capability to trigger alarms for security operations centers (SOCs) to investigate such activities further. In addition, it allows them to understand whether there are opportunities for the organization to scale back local users, groups and permissions based on activity within the environment.

Data Security

When moving data to a third party's infrastructure, data protection becomes a priority for organizations. CASBs afford key capabilities in a single tool: from understanding where your data resides, to determining which data is being transmitted back and forth, to uncovering object storage that does not offer malware detection. They integrate with cloud storage services such as Amazon Simple Storage Service (Amazon S3) and provide analysis of data as it is transferred to and from Amazon S3 storage.

Depending on how an organization integrates the CASB, the level of data security can vary. For example, CASBs may also offer integrations to existing endpoint data loss prevention (DLP) tools, such as McAfee DLP Endpoint and McAfee MVISION Cloud. This proxied connection accomplishes a very important task organizations must consider when moving assets to the cloud: DLP. Understanding what data is being transferred and where it is going becomes a unique challenge. Because of this, a CASB has an

² This paper mentions product names to provide real-life examples of how various tools can be used when integrating CASBs into cloud environments. The use of these examples is not an endorsement of any product.

opportunity to aid in enforcing DLP policies, as well as providing malware detection for data that is coming to and from the organization's AWS footprint or cloud applications. This is a desirable attribute of CASBs.

“With a CASB’s capability to perform UEBA, security personnel can get a better understanding of behavior that deviates from the norm.”

Configuration Compliance

With elasticity and self-service being a couple of the key characteristics of cloud computing, development staff need to be able to understand whether a workload’s configuration is adhering to a best practice. This is true both for those teams that are experienced and have been using AWS for several years and for new teams that are leveraging AWS for the first time.

A CASB can also provide:

- **Configuration reporting** — A CASB can extend itself by evaluating AWS accounts, looking at the configuration(s) and aligning them to best practices such as discovering shadow IT cloud services. The Center for Internet Security (CIS) Benchmarks³ help organizations identify best practices.
- **Compliance reporting** — To further the configuration reporting and aligning to best practices, a CASB can provide insight to the compliance status. Security personnel can also determine whether controls for SaaS applications are being enforced.

With a CASB monitoring configurations within their AWS footprint and cloud applications, security personnel can identify at-risk workloads and correct them, as well as gain understanding of additional applications that may be in use. A common at-risk configuration could be that multifactor authentication

(MFA) is not enabled on an AWS Management Console user. The CASB can display an alert. For application discovery, ask if data is going to a data storage service that is not on a preapproved list of cloud-based storage or if you detect that data is going to another third-party service.

Use Cases

Table 1 provides three use cases and the features security teams can leverage with a CASB to address identified needs. These use cases are common for a CASB, and solutions can be addressed by AWS Marketplace partners such as Netskope and McAfee.

CASBs are growing in popularity, as the uses of cloud applications and IaaS are increasing. Being able to integrate security solutions into these platforms is becoming necessary for organizations that consume these modern technologies.

Use Case	How to Leverage
I want to enforce DLP policies.	Using a CASB in the API configuration allows the organization to enforce DLP policies to control its data based on signatures that it creates. The organization is also able to leverage the reverse proxy or forward proxy configurations to inspect data in transit and block data transfer violations.
What cloud applications are being used in the organization?	Because a CASB is acting as a “middleman-as-a-service,” an organization is able to gain insight into other cloud applications. This helps the organization prevent shadow IT in the modern cloud world.
Are there unused identity and access management (IAM) groups in any of the organization’s AWS accounts?	A CASB gathers this data directly from the organization’s AWS account. It does not need to utilize the different deployment models.

Table 1. Use Cases

“When moving data to a third party’s infrastructure, data protection becomes a priority for organizations. CASBs afford key capabilities in a single tool.”

³ Center for Internet Security, www.cisecurity.org/benchmark/amazon_web_services/

Summary

Leveraging a CASB has several advantages for an organization. As an organization moves applications and data from an on-premises data center to the cloud, the number of applications that it can leverage grows constantly, as do the areas where data can reside. Strategizing how an organization is going to secure cloud-based infrastructure, applications and data are a few of the pieces to consider when moving to the cloud. When determining what type of CASB deployment to use, an organization should consider what kinds of devices are within the organization, whether managed, unmanaged or both. An organization also needs to consider if it has the capability to manually and automatically push agents or configuration files to the workstation. Once an organization evaluates the deployment method and functionality, it possesses the ability to maintain the deployment it selected.

About the Author

Kyle Dickinson teaches SANS SEC545: Cloud Security Architecture and Operations and has contributed to the creation of other SANS courses. He is a cloud security architect for one of the largest privately held companies in the United States. As a strategic consultant in his organization, Kyle partners with businesses in various industries to better understand security and risks associated with cloud services. He has held many roles in IT, ranging from systems administration to network engineering and from endpoint architecture to incident response and forensic analysis. Kyle enjoys sharing information from his experiences of successes and failures.

Chapter 14: How to Build a Security Visibility Strategy in the Cloud



Dave Shackleford

SANS Senior Instructor & Author

"With the shift to cloud, security and operations teams have had to look at different tools, processes and services that are geared toward software-defined infrastructure. For lots of reasons, developing visibility across all cloud assets and accounts is a top priority. On the one hand, there are more advanced methods and tools built into the cloud provider fabric to help gain a continuous monitoring view of the environment. On the other hand, there's more to cover, including the cloud control plane and a vast array of new services running in the cloud.

To improve visibility in the cloud, especially with the dynamic nature of today's deployment and runtime life cycles, organizations need to enable all the tools they have in the arsenal to help. This includes cloud-native logging, cloud integrated monitoring, workload introspection capabilities, and much more. At the same time, the SOC must update use cases and workflows to craft entirely new incident response playbooks, too. This chapter provides security analysts with the tools and concepts necessary to craft a more cloud-centric security visibility strategy."

Introduction

Today organizations are storing sensitive information ranging from business intelligence to personally identifiable information, health records, credit cards and other regulated data in the cloud. It is obvious that cloud is here to stay, and security professionals need to manage the threats and vulnerabilities that go along with cloud deployments. The good news is that more powerful tools and capabilities are available in the cloud than ever before, and this all starts with increasing visibility for cloud implementations, both with cloud-native tools and services and third-party tools and products that have been adapted to cloud provider environments. In this paper, we look at a variety of controls to ensure network, application, instance/ container, database/storage, and control plane visibility and build upon them to create a security visibility strategy for the cloud.

Types of Security Visibility Needed in the Cloud

The two major types of visibility that security teams need to focus on in the cloud today are:

- **Event-driven visibility** — The most common types of visibility that security teams have traditionally focused on are events. These events can be derived from a wide variety of sources, including operating system logs, application logs, network device and platform logs and events, and security system events (intrusion detection and prevention, data protection tools, anti-malware platforms and more). In the cloud, all of these events still have merit and all can (and should) be collected as needed. However, the cloud service environment itself can also track events occurring across infrastructure, so security teams have a new category of events they can use to monitor for unusual or suspicious activity. For example, a security operations center (SOC) can monitor AWS CloudTrail¹ events for an Amazon Elastic Compute Cloud (EC2) instance spawned from a non-approved machine image or a user attempting to deactivate multifactor authentication (MFA).
- **Behavior-driven visibility** — The other major types of visibility needed in many environments are more driven by events occurring over time, indicating a pattern of behavior. Particularly in cases of insider abuse, account hijacking and illicit use of cloud resources, organizations need insight into larger datasets over longer periods of time to really see whether unusual or malicious activities are afoot.

¹This paper mentions product names to provide real-life examples of how visibility tools can be used. The use of these examples is not an endorsement of any product.

With these two types of visibility in mind, the next section describes the types of controls you will need to ensure security visibility.

Security Visibility Today

The importance of visibility into what the environment looks like and the inventory of available assets cannot be overstated. The first of the Center for Internet Security (CIS) Critical Security Controls² focuses entirely on shoring up this lack of visibility through maintaining a sound inventory of systems operating within the environment. The security concept “You can’t secure what you don’t know about” holds true in any environment, and this control has been the highest-priority control since the list’s inception. The second CIS Critical Security Control focuses on gathering and maintaining an inventory of software running on systems. Both of these controls fit into the identify function of the NIST Cyber Security Framework (CSF), which is illustrated in Figure 1.

- **Network visibility** — The types of controls often used to achieve network visibility include network firewalls, network intrusion detection and prevention, load balancers, proxying tools, and network flow data (behavioral) collection and monitoring. Leading network vendors have adapted products in all of these categories to integrate into a virtual private cloud (VPC) architecture, granting network and security teams the same security capabilities and insight into network traffic they’ve attained internally. Cloud-native access controls such as security groups and flow logs enable security teams to monitor and track network events and behaviors.
- **Application visibility** — Application visibility relies on tracking events and behaviors at scale as workloads communicate within the cloud environment as a whole, in addition to the local application logs on individual systems and containers. Developing true application visibility often relies on feeding events into event management and SIEM platforms, which have also been well adapted into cloud environments, often via API integration.

“The importance of visibility into what the environment looks like and the inventory of available assets cannot be overstated.”

² www.cisecurity.org/controls

NIST Cyber Security Framework				
IDENTIFY	PROTECT	DETECT	RESPOND	RECOVER
Asset Management	Access Control	Anomalies and Events	Response Planning	Recovery Planning
Business Environment	Awareness and Training	Security Continuous Monitoring	Communications	Improvements
Governance	Data Security	Detection Processes	Analysis	Communications
Risk Assessment	Info Protection Processes and Procedures		Mitigation	
Risk Management Strategy	Maintenance		Improvements	
	Protective Technology			

Figure 1. The NIST Cyber Security Framework

- **Instance/container visibility** — Logs and events generated by services, applications and operating systems within cloud instances should be automatically collected and sent to a central collection platform. Automated and remote logging is something many security teams are already comfortable with, so organizations implementing robust cloud security designs really just need to ensure that they are collecting the appropriate logs, sending them to secure central logging services or cloud-based event management platforms, and monitoring them closely using SIEM and/or analytics tools. In the case of containers and container management tools, many new and well-known providers of vulnerability scanning and configuration assessment services have adapted their products to work in the cloud, granting deep visibility into both container image configuration and runtime event monitoring.
- **Database/storage visibility** — Many cloud deployments employ a wide variety of storage types, including block storage, blob-type storage, databases and more. Security visibility for storage components often revolves around access controls and permissions, as well as events related to encryption and other protective measures implemented within the storage platform. All major cloud storage types include various forms of logging, and many include access control measures. Many encryption and data monitoring tools are available for public cloud storage, as well.

- **Control plane visibility** — Another type of visibility that is now available in the cloud is of the cloud environment itself: the control plane. In addition to extensive logging of all activity within the environment itself, a number of new services are available to continuously monitor cloud accounts and environments for best practices configuration and security controls status. Imagine a single service to monitor the entire data center and its configuration all at once!

Myths About Cloud Security Visibility

As cloud adoption has increased, a couple of myths about cloud security visibility linger.

“We can’t get adequate logging in the cloud.”

Today, this statement is blatantly false, because major infrastructure-as-a-service (IaaS) providers have enabled extensive logging of all activity within the environment, essentially recording every API call made in any way.

“Network security visibility is less capable in the cloud.”

With the right mix of tools and architecture, this is also untrue. More and more, leading network security providers are adapting products to integrate into leading IaaS clouds, and coupled with cloud-native network controls, this provides plenty of opportunity to see and control traffic.

What is Different About Visibility in the Cloud?

One major development in cloud security that immediately benefits security teams is the reality that cloud-based assets are inextricably linked to the provider's environment, making them always visible. Through a combination of integrated APIs, scanning and local agents, it is possible to improve upon inventory and asset management strategies more than ever. In essence, there's an "always-on" level of visibility that teams can query and monitor, and there's really nowhere to hide in the cloud.

In addition, as noted earlier, a comprehensive control plane is now part of the mix for security-related tasks and operations. What does this mean to visibility? In essence, the environment (and APIs offered by the cloud provider) becomes a unified backplane that organizations can attach monitoring tools to, generate event data from, and set event and behavior "triggers" around that puts this control plane to work for security teams in an automated fashion. By building out policies for event monitoring, continuous scanning of workloads and events, and potentially responding through automated actions, the cloud platform lends itself to deeper levels of visibility than were possible in traditional data center environments. Imagine having a single control plane for your entire data center, where all tools could be connected, events generated and monitored, access managed and so on—this is truly what's possible in the cloud.

All of this is possible, of course, because the entire environment is software-defined. In addition to adapting existing tools and services to work within the new control environment, many services from the cloud providers themselves are emerging to augment security operations strategies. It is possible to have more than one tool or service monitoring various facets of cloud environments at all times—with minimal additional overhead.

Building a Cloud Security Visibility Strategy

The first function outlined in the NIST CSF is Identify, which consists primarily of asset management, governance and risk assessment practices and controls within the environment. Accordingly, the first step to building a cloud visibility strategy is to determine what types of event data and information are available in the cloud environment you're operating within, which can immediately help to achieve the goals of the identify phase. Aside from agent-based tools that can help to collect workload and container events, and other third-party platforms that organizations may choose to implement (discussed shortly), logs and events that contribute to cloud visibility also include environment logs that describe interesting API activity (which would also align under the investigate function of the NIST CSF). Take, for example, an

AWS CloudTrail event that indicates a cloud user trying to deactivate an MFA device, as shown in Figure 2. Be sure to evaluate these log types carefully to understand what types of information they provide you.

```
"eventTime": "2017-01-20T18:53:02Z",
  "eventSource":
    "iam.amazonaws.com", "eventName":
    "DeactivateMFADevice",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "1.2.3.4",
  "userAgent":
    "signin.amazonaws.com",
  "requestParameters": {
    "userName": "dave",
    "serialNumber": "arn:aws:iam::000012345678:mfa/dave"
  },
  "responseElements": null,
```

Figure 2. Suspicious AWS CloudTrail Event

Another major element of the NIST CSF is Protect, which emphasizes many security controls that would be involved in improving security visibility. Such controls include firewalls and security agents that can aid in protecting from malware, network behavior monitoring, event management tools and more. Consider the following process to select and implement the most effective cloud security visibility strategy:

1. Be sure to investigate third-party options from vendors and service providers that can enhance and augment your monitoring and visibility strategy.
2. Before considering the latest cloud-native tools and capabilities from cloud providers, consider the critical factors that may dictate when you should keep your in-house vendor products in place (or possibly choosing entirely different third-party tools versus those you've had) as opposed to moving to new cloud service provider offerings. Sticking with your current tools makes sense if:
 - You have a well-supported vendor product that has been adapted to the cloud and scales well.
 - You have a highly distributed cloud deployment and need to keep operational overhead and skills to a bare minimum.

- Your vendor product has clear and distinct advantages over the cloud provider services offered and these make a difference to you.

In some cases, however, a combination of both vendor and cloud provider services/controls may make more sense than one solution alone. To that end, be sure to evaluate cloud-native controls that the provider offers. In-house services may offer simpler operations, better performance, improved capabilities, or deeper and more natural integration than existing tools. For many large enterprises, though, cloud-native solutions will be better implemented to augment and enhance security visibility alongside third-party tools.

Finally, make sure you tie together event monitoring, vulnerability scanning/monitoring and control plane visibility to create a true continuous monitoring strategy.

Building a Cloud Security Visibility Strategy

What does a modern cloud-enabled SOC look like for hybrid architectures? Figure 3 illustrates key issues a cloud-aware SOC should be prepared to work through.



Figure 3. Planning Steps

Architecture Planning

The SOC team needs to align with cloud architecture and engineering teams that have built the hybrid architecture and maintain it. DevOps teams will also be involved in governance and oversight of cloud activity monitoring and visibility, because they will be responsible for application development and deployments into a platform-as-a-service (PaaS) or IaaS environment.

The SOC team should strive to understand the following with the assistance of these teams:

- **What connectivity does the public cloud provider have back to the data center or primary operations location?** In many hybrid architectures, this connection is either a point-to-point IPsec VPN tunnel (or several of them), a dedicated telecommunications circuit of some fixed bandwidth, or a combination of both. The means of connectivity will determine accessibility into the cloud network environment, as well as bandwidth constraints on event data and other visibility information the SOC needs.
- **Are the appropriate tools enabled?** Discuss whether any deployment tools in use for managing and promoting infrastructure as code (code repositories, deployment tools like Jenkins, or template formats like CloudFormation, Terraform, etc.) should be enabled for auditing activities and access logging.
- **How will deployment images and container builds be deployed?** Discuss deployment images and container builds, so that the SOC understands where and how these will be deployed. Team members need to understand topics including image update cycles, storage locations and workload lifecycle to better enable contextual monitoring.
- **What are our plans for elasticity and scaling?** Discuss any plans for elasticity and automatic scaling operations that could increase or decrease activity and operations in the cloud environment. SOC teams must understand these issues so that they can better prepare to monitor the events and track changes accordingly.

Enabling Security Controls

The SOC should then enable the following options in various security control categories to ensure visibility is maximized in the cloud:

OS Hardening and Logging

Enable auditing and logging of all instances and containers to be forwarded to a central in-cloud storage location, where the data can then be streamed to an on-premises or in-cloud SIEM. Ideally, CIS guidelines and other industry benchmarks are built into deployment templates and images, and additional logging and hardening scripts can be created by experience over time.

Control Plane Logging

Ensure that all cloud provider control plane logging (such as AWS CloudTrail) is enabled and that these logs are being centrally collected and streamed to an on-premises or in-cloud SIEM through API integration. Any third-party services performing independent control plane logging and monitoring should be generating events and logs that can ideally be extracted via API and centralized within a SIEM or analytics platform. In addition, enable cloud-native behavioral analytics tools to monitor account behavior and activity specifically.

Identity and Access Management (IAM)

All directory service logs should be centrally collected, as should other logs such as central policy coordination through tools like identity and access management tools offered by cloud providers. Because most IAM users and groups tend to be service accounts and unique DevOps, testing and administration accounts, be sure to carefully monitor all activity pertaining to these users and roles. Any addition, deletion or changes of IAM policies should be noted carefully and prioritized, too.

Endpoint Security

Ideally, SOC teams will have installed and enabled endpoint detection and response (EDR) agents from a trusted third party or leading open source project, including tools that perform host IDS functions. Send all these events to a monitoring console that can integrate with SIEM and analytics tools.

Network Security

A SOC team should enable next-generation firewall (NGFW) platforms that offer intrusion prevention and detection, along with traditional network protocol and service/ port control. Also, enable and send cloud DNS logs and network flow records to a central monitoring platform that can feed data to SIEM and analytics tools.

Vulnerabilities/Configuration

Set up a best-of-breed third-party network and application vulnerability scanner to feed vulnerability reporting data back to a SIEM or analytics platform, and use a cloud-native scanning tool (if available) to enable more continuous monitoring (if available). Any continuous monitoring tools that the cloud provider offers should also be enabled to scan for specific conditions. For example, are all running workloads being started from approved images?

Threat Detection

With the proper visibility in place through logging and monitoring, along with large-scale analytics and data processing tools and capabilities, cloud consumers can now track and monitor both control plane activity (covered earlier) and threats from both internal and external sources over time. With a more complete picture of behavior, organizations can detect malicious, suspicious, and accidental/unintended actions and events.

Adapting Existing Processes and Functions

Finally, a SOC needs to adapt some of its existing processes and functions to properly improve visibility into their deployment of hybrid architectures. Take the following example of a traditional SOC walkthrough (see Figure 4).




Figure 4. Process for Adapting Processes and Functions

Initial Event

Based on collection and large-scale analytics processing of flow logs within their SIEM, SOC staff is alerted to a workload in a cloud subnet scanning or trying to communicate with other subnet members. These are recorded as **REJECT** messages from a number of ports where the subnet attempted

communication. Simultaneously, a serverless function that autotags instances exhibiting these scanning behaviors is triggered, adding the tag **Suspicious** to the instance with a value of **Yes**.

Within the same time frame as this initial alert, additional correlating evidence appears implicating strange behavior patterns on the part of an IAM account used in application interactions with this same system. The account was invoked from a remote command-line installation versus internal-only invocation.

Event Validation

Using a dedicated account with specific programmatic access privileges into the production environment, the SOC team runs a query to find out the instance configuration details based on the image it was deployed from, as well as how long the instance has been running and its remote IP address (if it has a public interface). Another SOC account query looks for any and all systems with the **Suspicious** tag every 30 seconds to see if new systems are appearing in the same subnet.

Investigation

Based on the behaviors seen, the SOC team runs a vulnerability scan on the workload to see if any obvious misconfigurations are present, or whether known vulnerabilities are found that could be exploited. At this point, the team can declare a formal investigation, open a ticket and initiate follow-up response and forensics processes.

Summary

The cloud has a lot to offer in the way of security monitoring and visibility. Organizations have the ability to capably monitor for both event-driven and behavior-driven activity, and now they have a single environment they can query for all the cloud control plane visibility they could ask for. Some adaptation of monitoring and preventive/detective tools may be required. However, organizations have more options because of the variety of cloud-native and third-party controls and services available. It is possible to implement and monitor the entire spectrum of control areas, ranging from network controls, including firewalls and intrusion detection services, to endpoint protection and monitoring agents, to continuous vulnerability scanning. Given large-scale analytics processing and numerous options to enable, collect, store and transmit log and event data from cloud assets and environments, organizations can more readily analyze everything happening in segments of their hybrid cloud networks and correlate this data with internal event information generated from existing security tools (some of which may be covering both internal and public cloud space).

About the Author

Dave Shackleford, a SANS analyst, senior instructor, course author, GIAC technical director and member of the board of directors for the SANS Technology Institute, is the founder and principal consultant with Voodoo Security. He has consulted with hundreds of organizations in the areas of security, regulatory compliance, and network architecture and engineering. A VMware vExpert, Dave has extensive experience designing and configuring secure virtualized infrastructures. He previously worked as chief security officer for Configuresoft and CTO for the Center for Internet Security. Dave currently helps lead the Atlanta chapter of the Cloud Security Alliance.

Chapter 11: How to Improve Security Visibility and Detection/Response Operations in AWS



Dave Shackleford

SANS Senior Instructor & Author

"Enabling logging in the cloud is easier than ever, but then what? What kinds of event data should organizations gather to be most effective? What else do they need to build an effective monitoring strategy that can then facilitate effective investigations and response? These are all common questions I hear frequently from security operations teams, and there are many types of workload, network and cloud control plane events that they need to collect in the cloud. That's just the beginning! After collecting this data, teams need to prioritize some types of events, integrate with both cloud-native and third-party monitoring tools and services, and leverage automation tools and controls to improve detection and response in highly dynamic environments. This chapter lays out the controls and services organizations should consider, identifies event data considerations in AWS for monitoring and alerting, and gives you some ideas on security automation, as well."

The Need for Cloud Security Monitoring

Security teams have increasingly realized a need to focus on monitoring tools and tactics for cloud environments. We've seen many types of cloud security incidents in the past several years, ranging from external intrusion attempts to internal misconfiguration and accidental exposure. Fortunately, cloud service providers (CSPs) have worked hard to create better cloud-native controls and services, as well as to enable third-party solutions to integrate with the cloud fabric for improved visibility and control. Security teams need to work diligently to update security monitoring and response practices to better reflect cloud-based tools and use cases.

In general, security teams need to focus on two major types of event monitoring in the cloud:

- **Event-driven monitoring** — The most common types of monitoring security teams have traditionally focused on are event-based. Events can be monitored from a wide variety of sources, including operating system logs, application logs, network device and platform logs, and security systems (intrusion detection and prevention, data protection tools, anti-malware platforms and many more). In the cloud, all of these sources are still important, and security teams can—and should—collect them all. However, the cloud control plane can also generate and track events occurring across an organization's infrastructure, so security teams can use a new category of events to monitor for unusual or suspicious activity. For example, a security operations center (SOC) could monitor events for an EC2 instance spawned from a nonapproved machine image or a user attempting to deactivate multifactor authentication (MFA).
- **Behavior-driven monitoring** — The other major type of security monitoring needed in many environments is driven by events that occur over time and indicate a pattern or trend in behaviors. Many use cases coincide with this model of monitoring, including cases of insider abuse, credential hijacking and illicit use of cloud resources. To best monitor for behaviors, security teams need access to and insight from larger datasets over longer periods of time to see whether unusual or malicious activities are occurring. An example might be an unusual pattern of workloads trying to communicate to other workloads within a subnet, potentially indicating system compromise and attempted lateral movement. This may be noted by observing large datasets of flow logs aggregated and monitored by a network monitoring solution or event management platform.

Cloud security monitoring and response increasingly focus on automation. While not all cloud security processes should be completely automated, there are many innovative automation capabilities built into the cloud control plane that can significantly improve many security monitoring and operations practices.

Collectively, logging and event monitoring, as well as automation strategies and tools, can enable security teams to build a continuous monitoring strategy in the cloud. This consists of two core strategies:

- Baseline monitoring and logging for workloads and the cloud control plane
- Scanning within the cloud for behaviors, conditions and vulnerabilities

Enabling Cloud-Native Event logs and Event Management

To establish baseline monitoring, security teams should gather and process the following:

- Cloud control plane logs (such as AWS CloudTrail¹ logs)
- Workload OS/application logs
- Network flow logs for virtual private clouds (VPCs)

Security teams should also leverage automation for improved operational capabilities with services like AWS Lambda and AWS Config.

Cloud Control Plane Logs

The first, and perhaps most obvious, step security analysts need to take is to collect logs from all relevant CSP environments. At the same time, analysts need to ensure that all the logs are going to a common location. An example of a cloud control plane logging service is AWS CloudTrail, which records any API calls made to Amazon Web Services (AWS). The service captures an extensive amount of data that security professionals will want to see, including the identity of the API caller, the time of the API call, the source IP address of the API caller, the request parameters and the response elements returned by AWS. AWS CloudTrail logging captures all requests made from the standard AWS Management Console, command-line tools, any AWS Software Development Kits (SDKs) and other AWS services.

¹ This paper mentions solution names to provide real-life examples of how cloud security tools can be used. The use of these examples is not an endorsement of any solution.

AWS CloudTrail solves one of the most challenging issues many security teams face when migrating IT resources into AWS: the capture and maintenance of cloud service event data that can feed log management and SIEM platforms. AWS CloudTrail uses Amazon S3 buckets for storage of the log data, allowing security teams to leverage the same APIs to access data quickly and easily for correlation and aggregation internally. Log data can also be automatically deleted after a certain period of time, or archived to internal storage or additional Amazon services like Amazon S3 Glacier for longer-term retention. Aggregation of log data across accounts and regions is possible, as is automated alerting and notification when certain events are registered. AWS CloudTrail log file integrity can also be enabled to hash all logs upon delivery and then monitor them afterward as well.

Most major CSPs allow logs to be downloaded from their environment (e.g., leading SaaS providers) or stored in a dedicated storage node (e.g., a dedicated S3 bucket). There are also a number of third-party security event aggregation and analysis platforms available for the cloud, including Sumo Logic² and others. These services may offer teams a simpler way to aggregate logs from multiple cloud services, and they often integrate more readily with these services through provider APIs.

Workload Security Events

The second type of logs that teams need to collect are those associated with different server and container workloads. You should collect logs from your instance OS, just as you would in your own data center. This means syslog, Windows events and all the other logs you'd normally try to collect for security and operational reasons. The basic mechanics of generating logs and sending them somewhere might be the same, in general, depending on the deployment model you have. Really, you should monitor these logs just like logs from your in-house systems. However, because of volume and cost, sending them to an in-cloud log collector and/or event management platform likely makes sense. This process is distinct from logging within the CSP environment, where you focus on API calls and access to the admin console for your cloud environment. It's important to make the distinction between cloud system monitoring and cloud environment monitoring. To ensure security, you must log and monitor systems just as you always have.

To enable consistent workload monitoring and logging, many organizations need to create and enable a central cloud log repository to store logs generated within workloads. There are many ways to accomplish this, but AWS has a unique agent, Amazon CloudWatch, that can be installed into Amazon EC2 workloads. This agent forwards syslog and other standard events to a dedicated Amazon CloudWatch logging group. From there, these logs can be parsed and analyzed, or streamed to a

² www.cisecurity.org/controls

different event management and monitoring solution through streaming services like Amazon Kinesis Data Firehose. For most organizations, the data export costs associated with large volumes of workload logs can prove somewhat prohibitive to simply sending all logs back to on-premises data collectors and SIEM tools. While this may work with a small volume of cloud services and workloads, large organizations will eventually want to enable cloud-native log collection and analysis tools instead.

“It’s important to make the distinction between cloud system monitoring and cloud environment monitoring. You must log and monitor systems just as you always have.”

Network Flow Logs

Another critical type of data that should be collected and monitored in cloud environments is network flow data. For all major clouds, this can be enabled at the virtual private cloud (VPC) level, and these flow logs can then be sent to a dedicated storage node for analysis. With AWS VPC Flow Monitoring, network and security teams can add network behavioral monitoring to their overall capability set, and these logs have a wealth of data that can prove useful in detecting strange patterns of access and behavior in the AWS environment.

Most network traffic is recorded in AWS, except for:

- Traffic between EC2 instances and Amazon DNS services
- Amazon Windows license activation traffic for Windows EC2 instances
- Multiple IP addresses traffic (only primary address is logged)
- Instance metadata traffic to and from **169.254.169.254**
- DHCP traffic

Analysts can use this data to detect unusual patterns of communication between instances and workloads in the VPC environment, as well as specific malicious or suspicious activities originating outside the cloud and targeting assets (for example, SSH brute-force attempts). Keep in mind that enabling this type of logging can produce a staggering quantity of event data, and you will need to leverage some sort of toolkit (SIEM, security analytics, etc.) to build behavioral baselines for monitoring purposes.

Improving Visibility in the Cloud

To improve security visibility in the cloud, security operations teams will want to develop a continuous monitoring strategy that uses a combination of cloud-native services and third-party options. This strategy provides the most comprehensive range of coverage for both proactively assessing the environment and detecting unusual events or anomalous behavior rapidly. Within AWS, for example, a continuous monitoring framework might include such services as:

- **Event-driven monitoring** — This service performs vulnerability assessments of your cloud instances. An agent is required to perform scans, and most operating systems are supported (at least most Linux and Windows OSes). Amazon Inspector provides a number of rules templates, including CVE (for listing missing patches and other typical vulnerabilities that a vulnerability scanner would report on), CIS Benchmarks (for industry-standard configuration and control practices), general security best practices and so on. Scans can run between 15 minutes and 24 hours. Longer scans are more thorough and provide better baselines. Longer scans can really help to evaluate state over time and may help you to detect the state of systems in a rapidly changing DevOps environment. Amazon Simple Notification Service (SNS) notifications can be queued to alert you or feed to scripts and automation engines like AWS Lambda.
- **AWS Config** — This configuration monitoring toolkit for your AWS systems can define your baseline image, monitor systems continually and alert whenever a system's configuration changes. AWS Config is natively integrated into AWS, and it can easily be set up to help keep your system state secure. Another key feature of AWS Config is its inventory capability. One advantage of the cloud is that nothing can hide, because all assets are 1) software-defined and 2) linked inextricably to the CSP's backplane. For this reason, the discovery and inventory elements of change and configuration management should be easier than ever! In the case of AWS Config, it doesn't get much easier—the service just finds everything and then lets you query

AWS to see what you have. Recent additions to the AWS Config service allow for automated remediation and alerting as well.

- **Amazon CloudWatch** — This service allows you to monitor data and events and create alarms based on events in your AWS environment. Amazon CloudWatch, which integrates with almost all AWS services, can collect and track metrics, monitor log files, initiate alarms and automatically respond to changes in your AWS environment. For this reason, it's one of the most flexible monitoring tools you can use.
- **AWS Security Hub** — This service offers basic continuous monitoring for AWS accounts, looking at CIS Benchmarks configuration checks and more. Additionally, a number of third-party security tools can integrate into AWS Security Hub to create a centralized dashboard of events and security monitoring and operations..
- **Amazon GuardDuty** — This service analyzes a vast volume of log and intelligence data (both internal to AWS and from third parties) to deliver threat intelligence about customer account behavior. Results from Amazon GuardDuty can be integrated into Amazon CloudWatch and other event-triggering systems in AWS, or sent to the SOC or other locations for analysis with different tools.
- **Amazon Detective** — This service collects and aggregates logs across AWS resources and performs deep analysis on them to detect behavior anomalies and other events for faster and more efficient root-cause analysis and investigations. This feature is still in preview as of early 2020.

Many organizations may want to integrate all cloud-based events—both workload events and cloud control plane events—into an existing centralized detection and response capacity (usually focused on integrating SIEM and other large-scale correlation platforms for cloud monitoring). There are cloud-integrated API connectors for all major SIEMs today, such as Sumo Logic, Securonix, Sonrai Security and more.

While this option is certainly a possibility, the costs to aggregate and export data (even over dedicated network connections like AWS Direct Connect) may be significant. For this reason, many organizations are now considering or implementing cloud-native SIEM tools.

“A continuous monitoring strategy that uses a combination of cloud-native services and third-party options provides the most complete range of coverage for both proactively assessing the environment and detecting unusual events or anomalous behavior rapidly.”

What to Look For: Enabling the SOC

Once cloud logs are being collected and aggregated, analysts need to sift through all the various events and start prioritizing them. There are several keys to this process, including:

- **Adding context** — If logs can be “tagged” as originating from a specific ISP or CSP, that can help provide context on the use cases of the service. For example, logs from identity management services like AWS Identity and Access Management (IAM) have a specific user context, whereas events from Amazon EC2 may need additional details about workloads to provide the proper context for evaluation.
- **Defining priorities** — Security analysts focused on the cloud must first decide what events and behaviors are most critical to monitor. Common starting points include any login activity to cloud management consoles; any changes or attempted changes to important cloud objects and data; any creation, deletion or modification of credentials or cryptographic keys; and attempts to modify or delete audit logs.
- **Tuning alerts** — Tuning is incredibly important for cloud logging and event management. You want to suppress redundant alerts, both those that are entirely operational in nature and those not directly related to security. To build appropriate behavioral baselines of events in the environment, you also likely need to allow several weeks or even months of data to accumulate. Make tuning a regular part of your weekly monitoring processes.

- **Housekeeping of accounts and credentials** — Leftover user credentials, cloud accounts and data can lead to potential risks in the cloud. Work closely with human resources teams to disable credentials to cloud accounts quickly, and monitor for all attempts to log in with disabled or deleted credentials for at least several weeks after a user has left the organization. It's a good practice to monitor user account activity of employees who have given notice to ensure that they don't try to take or sabotage critical data. For example, look for sudden increases in data exports, transfer or overall account use.

Another area of focus for cloud events should be the originating point of cloud activity. Security teams should consider a login from a new country or location where the organization doesn't do business or have users to be a very high priority alert. Many cloud logs include enough detail to note where the login came from.

Identification and Prioritization of Potential Events

Where to start? Security operations teams might feel somewhat overwhelmed when starting to sift through cloud logs and events. Fortunately, many types of events and information can help identify potential incidents in the cloud, including:

- **Incident notifications from your CSP** — This depends on your CSP model and deployment type, as well as contractual SLAs and terms.
- **Billing alarms** — These are key! If you have a reasonable idea of a monthly billing range, you can break this down to define "checkpoints" of what your bill should be at any given time. If these thresholds are crossed, a billing alarm could alert you and investigate what is causing the additional cost
- **IAM activity (logins in particular)** — Monitor your user activity within the cloud. In particular, monitor admins carefully, because these user credentials are prime targets for attackers. Any nonfederated user access should also be a high priority.
- **Cloud environment logs (e.g., AWS CloudTrail)** — General API logs can tell you when instances are created or changed, when storage attributes change and so on. Focus on the types of events that could be problematic to the environment. These event types include access or changes to critical assets, modification of identity policies, deletion or changes to cryptographic keys, and so on.

As a general rule, security operations teams should prioritize the following types of events (listed by order of priority/severity):

Priority 1

- Launching a workload that is not from an approved template
- Launching any containers from unapproved images in a repository
- Launching any assets in unapproved regions
- Modifying any IAM roles or policies
- Modifying or disabling cloud control plane logging or other security controls
- Logins to the web console (unauthorized)

Priority 2

- Unusual user behaviors (trying to access unauthorized resources, etc.)
- Adding/updating new workload images
- Adding/updating new container images
- Logins to the web console (authorized)
- Updating/changing serverless configuration

Priority 3

- Changes to security groups or network access control lists (ACLs)
- Updating/changing serverless function code

Identification and Prioritization of Potential Events

A cloud monitoring workflow should ideally look like one shown in Figure 1.



Figure 1. Cloud Monitoring Workflow

Logging begins with a central logging engine like AWS CloudTrail and/or a log collection agent from a SIEM solution extracting log data from a data store (primarily for workload logs if applicable). All logs, irrespective of source, need to be monitored for suspicious activity in the context of what environment the assets operate within, with detection filters set up to send alerts or perform more automated response actions. Any security operations team should spend time with all cloud environment logs to better understand the behavior of the workloads and services operating there.

For example, AWS CloudTrail captures an enormous range of event data, and tools like Amazon CloudWatch enable you to search for many different events. Table 1 lists some examples of starting points.

Table 1. Starting Points for Event Searches	
AWS CloudTrail Event	Reason for Investigation
ConsoleLogin	A user initiates console login activity.
StopLogging	A user tries to stop AWS CloudTrail.
CreateNetworkAclEntry	Someone creates a network ACL, which could expose attack surfaces or vectors.
CreateRoute	Someone creates a new route for data path control, which could expose attack surfaces or vectors.
AuthorizeSecurityGroupEgress AuthorizeSecurityGroupIngress RevokeSecurityGroupEgress RevokeSecurityGroupIngress	Monitor all changes to security groups.
ApplySecurityGroupsToLoadBalancer SetSecurityGroups	Security group changes that tie to elastic load balancers are interesting, often in scaling operations. This may indicate unusual traffic surges in the environment.
AuthorizeDBSecurityGroupIngress CreateDBSecurityGroup DeleteDBSecurityGroup RevokeDBSecurityGroupIngress	Amazon RDS instances have a different nomenclature for security groups, but are the same thing conceptually. Security teams should monitor such instances.

Additionally, there are a number of serverless events in AWS Lambda that could prove to be interesting starting points. For example, if someone deletes a function (**DeleteFunction**), this might be important. The same could apply for RemovePermission.

Table 2 lists the most critical AWS Lambda events to monitor immediately for security.

Security teams also need to be proactive in securing the cloud environment. Security operations and engineering teams should work with cloud operations and engineering teams to implement more effective controls around:

- **IAM and privileges (and credential security)** — This can be one of the most difficult areas to solidify in cloud security, because there are many types of privileges and roles that can be defined. AWS has a service called AWS IAM Access Analyzer, which is free and integrated into the AWS IAM platform. This service can help with assessing any AWS native or custom IAM policies to determine where excessive or unintended privilege allocation may be present based on AWS best practices and assigned users/groups.
- **Resources and resource utilization** — Cloud control plane logs from services like AWS CloudTrail can (and should) be heavily leveraged to monitor new, modified and deleted assets in the environment, as well as access to assets and service interaction in the cloud environment. These logs need to be integrated with a SIEM and/or cloud-native cloud monitoring solution like Amazon CloudWatch to build the appropriate triggers for alerting, as well as monitoring and reporting metrics as warranted. Some behavioral trending over time can also be assessed and reported through analytics tools like AWS Security Hub and Amazon GuardDuty, as well.
- **Activity in specific regions** — One of the best quick wins for security teams is to purposefully disable all geographic regions not in use; a follow-up to this is enabling explicit monitoring for cloud control plane logs (like AWS CloudTrail) to look for any activity in regions marked as “not in use” or “disabled.” A common tactic intruders use for malicious activities like cryptocurrency mining is to create unauthorized assets and workloads in unused regions to “buy time” before detection. Teams should consider any alert for activity in an unauthorized or unused region a high priority.

Regardless of the tools chosen, SOC teams need to adapt their workflows and monitoring processes to include as much log and event data from the cloud as possible. This invariably requires significant effort to better learn and understand the patterns of events and service interaction in the cloud environments

chosen. Spending some time each month or quarter developing “game day” or tabletop exercises to flesh out cloud monitoring and response use cases is an excellent way to engage the SOC team in cloud initiatives and improve the team’s skills and processes at the same time.

Table 2. Events for Immediate Monitoring	
AWS Lambda Event	Reason for Monitoring
DeleteEventSourceMapping	Someone could delete the data source that triggers an AWS Lambda function, making it “blind.”
DeleteFunction	A function could be deleted purposefully or accidentally, leading to security issues.
RemovePermission	This could lead to a lockout scenario or lack of access when needed (think IAM service account or role access to AWS Lambda).
UpdateEventSourceMapping	Data could be pulled from a different source, leading to incorrect function results.
UpdateFunctionCode	The function could be broken or tampered with to prevent security-specific functionality from executing (for example, by adding comments).
UpdateFunctionConfiguration	The configuration of the function could be changed to limit its resources, causing poor or flawed execution.

SOAR and the Role of Automation

Increasingly, more enterprise incident response teams are actively looking for opportunities to automate processes that often take up too much of their highly skilled analysts’ time, as well as those processes that require lots of repetition (and may provide little value in investigations). Common activities that many teams consider for automation include the following:

- **Identifying and correlating alerts** — Many analysts spend inordinate amounts of time wading through repetitive alerts and alarms from many log and event sources, and spend time piecing together correlation strategies for similar events. While this is valuable for later stages of investigations, it can also be highly repetitive and is therefore a good candidate for some degree of automation.
- **Identifying and suppressing false positives** — This work can be tedious on a good day, and overwhelming on a bad one. Identifying false positives can often be streamlined or automated using modern event management and incident response automation tools.

- **Initial investigation and threat hunting** — Analysts need to quickly find evidence of compromise or unusual activity, and often need to do so at scale.
- **Opening and updating incident tickets/cases** — Due to improved integration with ticketing systems, event management and monitoring tools used by response teams can often generate tickets to the right team members and update these as evidence comes in.
- **Producing reports and metrics** — Once evidence has been collected and cases are underway or resolved, generating reports and metrics can take a lot of analysts' time.

Examples of security response automation include:

- Automated DNS lookups of domain names never seen before
- Automated searches for detected indicators of compromise
- Automated forensic imaging of disk and memory from a suspect system, driven by alerts triggered in network- and host-based anti-malware platforms and tools
- Network access controls automatically blocking outbound command and control (C2) channels from a suspected system

A fair number of vendors and tools can help integrate automation activities and unify disparate tools and platforms in use for detection and response. These include Swimlane, IBM Resilient Incident Response Platform³ and more, most of which leverage APIs with other platforms and tools to allow them to share data and create streamlined response workflows. Factors to consider when evaluating these automation tools include maturity of the vendor, integration partners, alignment with SIEM and event management, and ease of use and implementation.

Incident response (IR) in the cloud may rely on scripting, automation and continuous monitoring more heavily than in-house IR currently does. Many of the detection and response tools emerging for the cloud are heavily geared toward automation capabilities. To effectively implement automated IR in the cloud, IR teams need to build automated “triggers” for event types that run all the time (such as

³ Swimlane is a trademark of Swimlane LLC; IBM and IBM Resilient Incident Response Platform are registered trademarks of International Business Machines Corp.

Amazon CloudWatch filters), especially as the environment gets more dynamic. Deciding what triggers to implement and what actions to take is really the most time-consuming aspect of building a semi-automated or automated response framework in the cloud. Do you focus on user actions? Specific events generated by instances or storage objects? Failure events? Spending time learning about cloud environment behaviors and working to better understand “normal” patterns of use is invaluable here.

“Factors to consider when evaluating security response automation tools include maturity of the vendor, integration partners, alignment with SIEM and event management, and ease of use and implementation.”

The following list provides a breakdown of the security automation model to consider for cloud deployments—it’s really broken into three major components:

- **Phase 1: Learn** — In this phase, you monitor for events occurring in the environment. With AWS, this would likely come from AWS CloudTrail logs, Amazon VPC Flow Logs, Amazon CloudWatch Logs, etc.
- **Phase 2: Trigger** — Based on some pattern matching, using Amazon CloudWatch alerts or even a SIEM like Sumo Logic, you then trigger some sort of follow-up action.
- **Phase 3: React/Respond** — The final phase is the actual action triggered during the automation. This could be an AWS Lambda function that performs an action, a vulnerability scan or an alert sent via SNS or other method.

The use cases for phases 2 and 3, where certain events trigger responses, vary widely. These might include tagging assets that are behaving suspiciously, disabling access keys or user/service credentials, changing a security group to one that is a “quarantine” zone without internet access, or simply alerting a group of SOC analysts. Security teams need to spend some time developing these automation use cases and then look into the tooling needed to accomplish these goals through cloud-native and third-party services.

Conclusion

The cloud has a lot to offer in the way of security monitoring and visibility. Security teams have the ability to capably monitor for both event-driven and behavior-driven activity, and they now have a single environment they can query for all the cloud control plane visibility they could want. Security teams need to adapt monitoring and preventive/detection tools in some cases, although they might have more options due to cloud-native and third-party controls and services that are rapidly expanding. Teams can implement and monitor the entire spectrum of control areas, too, ranging from network controls like firewalls and intrusion detection services to endpoint protection and monitoring agents to vulnerability scanning continuously. With large-scale analytics processing and numerous options to enable, collect, store and transmit log and event data from their cloud assets and environment, teams can more readily analyze everything happening in this part of the hybrid cloud network and correlate this data with internal event information generated from existing security tools (some of which may be covering both internal and public cloud space).

That said, there's still a lot of work for SOC teams to do in reviewing events and building detection and response use cases. Building effective correlation cases for cloud monitoring can also be readily accomplished with the tools and services available today, but it will take time and a better understanding for SOC teams to adapt to different event sources and types.

One area of significant promise is automation—teams have all the event details they need, as well as tools and services to store and process them. With SOAR solutions and cloud-native processing and automation engines, security operations teams should see definitive improvements in their detection and response capabilities, because the cloud is a unified fabric with innumerable APIs to employ (for querying information and for performing detection, response and mitigation). As infrastructure becomes progressively more software-defined, this will be more and more important to security professionals everywhere.

About the Author

Dave Shackleford, a SANS analyst, senior instructor, course author, GIAC technical director and member of the board of directors for the SANS Technology Institute, is the founder and principal consultant with Voodoo Security. He has consulted with hundreds of organizations in the areas of security, regulatory compliance, and network architecture and engineering. A VMware vExpert, Dave has extensive experience designing and configuring secure virtualized infrastructures. He previously worked as chief security officer for Configuresoft and CTO for the Center for Internet Security. Dave currently helps lead the Atlanta chapter of the Cloud Security Alliance.

<https://t.me/learningnets>

Chapter 16: How to Build a Threat Detection Strategy in AWS



David Szili

SANS Certified Instructor

"AWS offers a plethora of data sources and services for security monitoring. The key to a successful threat detection program is to pay attention to instances and images, as well as to cloud network infrastructure and cloud management.

In this chapter, I discuss the main strategic steps, starting with the most critical data sources available such as Amazon VPC Flow Logs or AWS CloudTrail. I address how to leverage traffic mirroring technology and use intrusion detection systems to alert on malicious activities. Finally, I describe a few security monitoring best practices and automation options when it comes to responding to incidents."

Introduction

One major concern security teams have is losing visibility and detection capabilities when their organization moves to a cloud. While this might have been true in the early days of cloud services, these days providers are announcing new threat detection features and offerings almost every month. These new services open up the possibility of adjusting traditional network- and host-based monitoring to support intrusion detection in the cloud.

In this paper, we focus on the key steps illustrated in Figure 1 to detect threats in Amazon Web Services (AWS) and gradually build a security monitoring strategy.

Threat detection and continuous security monitoring in cloud environments have to integrate security monitoring of instances and images (system monitoring), just as they do on premises. For cloud services, however, it is also crucial to include the monitoring of the cloud network infrastructure and cloud management plane (cloud monitoring).

In terms of system monitoring, organizations must collect system logs and vulnerability scan results. They must also check the integrity and compliance of instances against policies and security baselines. The collection of operating system logs can be challenging because they require centralized collection for analysis and correlation.

Given the volume of this data and the associated cost of sending it back to an on-premises solution, using an in-cloud log collector or event management platform can be a much more viable option.



Figure 1. Steps to Build a Security Monitoring Strategy

As for the AWS Cloud environment, security teams must monitor admin access, changes made to the environment, API calls, storage and database access, and any access to sensitive and critical components. In the following sections, we explore data sources and services that help with event management and analysis.

The focal point of the threat detection strategy is to collect data from systems, networks and the cloud environment in a central platform for analysis and alerting. AWS Security Hub¹ is a service that automates the collection process and organizes and prioritizes security alerts into a single, comprehensive view. The data sources, services and solutions described in this paper all feed into this monitoring solution to provide visibility and detect threats.

Data Collection

The first step in creating a security monitoring strategy is to identify the available data sources and determine how to collect data from them. Key data sources include endpoint detection and response (EDR) tools, flow logs, data from intrusion detection and prevention tools, and alerts from Amazon GuardDuty (discussed in the “Event Management and Analysis” section) and other AWS tools. When considering data collection for security monitoring, the winning strategy is to focus on the data sources with the highest value and the best cost–benefit ratio—and to do so efficiently. AWS Security Hub simplifies data collection from a variety of sources and collects alerts into a single, comprehensive view, as described in the “Event Management and Analysis” section.

In the case of AWS, these are Amazon VPC Flow Logs and AWS CloudTrail logs. Amazon VPC Flow Logs provide visibility into VPC and instances network traffic. Flow records are small and have a fixed size, making them highly scalable, with longer retention times, even for large organizations. AWS CloudTrail provides the logs for monitoring the AWS Cloud environment itself. We examine these two data sources next.

“Focus on the data sources with the highest value and the best cost–benefit ratio—and do so efficiently.”

Flow Logs

Flow records, such as NetFlow or IPFIX, are a statistical summary of the traffic observed. Common attributes allow grouping of packets into a flow record. These attributes are the source and destination IP addresses, the source and destination ports, and the network protocol (usually TCP, UDP or ICMP). As a result of this summary nature of the flow records, they do not contain information about the application layer. Therefore, visibility is limited to Layer 4 and below. Flow logs still offer means to:

- Scope a compromise and identify communication with known attacker addresses.
- Identify large flow spikes that might suggest data exfiltration.
- Identify large counts of frequent, small traffic bursts that may be command and control traffic.
- Detect strange patterns of access and behavior.

Because a significant portion of today's network traffic is encrypted and application data is unavailable for analysts, the lack of Layer 7 information in flow records is of little concern. Flow analysis techniques work exactly the same for both encrypted and unencrypted communications. This makes flow analysis a great method for threat hunting without the need for SSL/TLS interception and full-packet capture.

The Amazon VPC Flow Logs feature enables security analysts to capture information about the IP traffic going to and from network interfaces in the VPC. Flow logs can be sent to Amazon CloudWatch or Amazon S3 buckets. A new log stream is created for each monitored network interface.

Amazon VPC Flow Logs records are space-separated strings. Similar to other flow records, such as NetFlow or IPFIX, they contain the network interface name, source and destination IP addresses and ports, number of packets, number of bytes, and the start and end times of the traffic flow.

One significant difference is that the flow record contains information on whether the security groups or network access controls lists (NACLs) permitted or rejected the traffic. The list of fields are as follows:

```
<version> <account-id> <interface-id> <srcaddr> <dstaddr> <srcport> <dstport> <protocol> <packets>  
<bytes> <start> <end> <action> <log-status>
```

The following flow record example is for NTP traffic (destination port 123, UDP protocol) that was allowed:

```
2 123456789010 eni-abc123deabc123def 172.31.32.81 172.31.16.139 59808 123 17 1 76 1563100613
1563100667 ACCEPT OK
```

This flow record example is for RDP traffic (destination port 3389, TCP protocol), which was rejected:

```
2 123456789010 eni-abc123deabc123def 172.31.9.69 172.31.32.81 44844 3389 6 20 4249 1563100613
1563100667 REJECT OK
```

Because VPC Flow Logs can produce a large quantity of event data, you will likely need a tool, such as a log aggregator and analytics platform or a SIEM solution, for monitoring and analysis (see the next section). For example, Amazon CloudWatch has a simple interface to search in log group events, but also has Amazon CloudWatch Logs Insights, which provides a powerful, purpose-built query language that can be used to search and analyze your logs. It is ideal for threat hunting and allows security analysts to use the techniques mentioned previously.

Amazon CloudWatch Log Insights has prebuilt sample queries for VPC flow logs, making it easy to get familiar with the query language and perform the analysis. These sample queries include cases like:

- Average, minimum and maximum byte transfers by source and destination IP addresses
- Top 10 byte transfers by source and destination IP addresses
- Top 20 source IP addresses with the highest number of rejected requests Security analysts must be aware that Amazon VPC Flow Logs exclude certain IP traffic such as Amazon DNS activity, DHCP or license activation

This is usually desired to avoid the duplication of information, for example, in the case of VPC mirrored traffic. In other cases, additional AWS solutions can fill in these gaps. For example, Amazon GuardDuty also monitors DNS traffic. Amazon VPC Flow Logs is an essential tool to leverage and should be collected in every VPC that has important assets.

API and Account Activity Logs

Cloud security also requires detailed visibility into user and resource activity. Actions that take place through the AWS Management Console, command-line tools or API services are just as important for preserving the integrity of cloud environments as they are for monitoring network activity and hunting for threats. This kind of event history helps in troubleshooting, change tracking and security analysis. The events should contain detailed information, including but not limited to:

- Time of the API call
- Identity of the API caller
- Source IP address of the API caller
- Request and response parameters

One of the first major additions to Amazon's security services was AWS CloudTrail, an AWS logging service that provides a history of any AWS API calls across accounts and Regions. AWS CloudTrail is enabled on your AWS account when you create it. From the AWS CloudTrail console, you can view, filter and download the most recent 90 days of events in CSV or JSON formats. You can also see the resources referenced by an event and pivot to AWS Config to view the resource timeline.

You can configure AWS CloudTrail trails to log management events and data events. Management events provide insight into management operations that are performed on resources in your AWS account. Examples include configuring security policies, registering devices and setting up logging. You can choose to log read-only, write-only, all, or no management events. Data events provide insight into the resource operations performed on or within a resource—for example, Amazon S3 object-level API activity or AWS Lambda function execution activity. To determine whether an AWS CloudTrail log file was modified, deleted or unchanged after it was delivered, you can enable log file validation.

AWS CloudTrail typically delivers log files within 15 minutes of account activity, and it publishes log files multiple times an hour, about every five minutes. The events are in JSON format, which makes them

humanly readable and easy to parse programmatically. The log entry in Figure 2 shows that a root user ("userIdentity": { "type": "Root") successfully signed into the AWS Management Console ("eventName": "ConsoleLogin") using multifactor authentication ("MFAUsed": "Yes"):

The event history feature allows you to perform simple queries and filter events in many ways, except for wildcard searches. You can use Amazon Athena for more in-depth analysis using standard SQL to interactively query the AWS CloudTrail log files delivered to the Amazon S3 bucket for that trail.

```
{
  "eventVersion": "1.05",
  "userIdentity": {
    "type": "Root",
    "principalId": "123456789010",
    "arn": "arn:aws:iam::123456789010:root",
    "accountId": "123456789010",
    "accessKeyId": ""
  },
  "eventTime": "2019-07-01T10:48:13Z",
  "eventSource": "signin.amazonaws.com",
  "eventName": "ConsoleLogin",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "1.2.3.4",
  "userAgent": "Mozilla/5.0 (Windows NT 10.0; Win64; x64; rv:68.0) Gecko/20100101 Firefox/68.0",
  "requestParameters": null,
  "responseElements": {
    "ConsoleLogin": "Success"
  },
  "additionalEventData": {
    "LoginTo": "https://console.aws.amazon.com/console/home?state=hashArgs%23&isauthcode=true",
    "MobileVersion": "No",
    "MFAUsed": "Yes"
  },
  "eventID": "3fcfb582-bc34-4c39-b021-10a394ab61cb",
  "eventType": "AwsConsoleSignIn",
  "recipientAccountId": "123456789010"
}
```

Figure 2. AWS CloudTrail Event Example

For an ongoing record of activity and events in AWS accounts, you have to create a trail and send events to an Amazon S3 bucket or Amazon CloudWatch Logs. Log data can be automatically deleted, or it can be archived to long-term storage, for example, in Amazon S3 Glacier. AWS CloudTrail provides exceptionally detailed visibility for AWS account activity, which is a key aspect of security and operational monitoring best practices.

Intrusion Detection and Prevention Systems

The second step in creating a security monitoring strategy is to determine how IDS/IPS fit into that strategy. Such systems have the same objectives in the cloud as on premises, such as alerting based on signature matching, behavioral anomalies and protocol mismatch. However, these solutions differ from the ones we have on premises, and because they must be adapted to the cloud environment, they might look less familiar at first. In a cloud environment such as AWS, you have control over your virtual machine instances and to your VPCs at some level, but not the physical network or the hypervisor platform (which includes components like virtual switches). The cloud service provider controls these lower layers; therefore, monitoring tools have to leverage the features provided by the upper layers.

Network IDS/IPS

On-premises network IDS/IPS (NIDS/NIPS) differs somewhat from cloud deployments. However, AWS offers additional features that enable network security monitoring. Hardware network taps or mirror ports (also known as SPAN ports) from hardware and virtual switches are not feasible because of the lack of Layer 2 access, but similar alternatives are available using agents or traffic mirroring. Security appliances that can be deployed in-line for monitoring or blocking can also be implemented in AWS.

One option is to send back all the traffic to on-premises sensors via a dedicated connection like AWS Direct Connect or through a VPN. This allows you to see traffic coming in to and out of the VPC, although on-premises sensors cannot see instance-to-instance traffic. Nonetheless, this model can be combined with the methods mentioned below for better coverage.

The other option is a do-it-yourself approach: using NAT instances or multihomed instances with multiple elastic network interfaces (ENIs) that can act as gateways and inspect traffic passing through them. This option results in more complex network design, extra configuration steps like the installation of NIDS/NIPS software or Linux traffic bridging, and additional resources to manage the platform, because there is usually no official support. Different instance types have a maximum number of network interfaces, and smaller instances typically only allow two.

A great alternative to the preceding approach is to use AWS Partner Network (APN) solutions from AWS Marketplace, which has major vendors like F5 Networks, Palo Alto Networks, Sophos and Check Point Software Technologies. Most NIDS/NIPS features could be handled by unified threat management (UTM) and next-generation firewall (NGFW) appliances from firewall vendors. These virtual appliances are also deployed in-line as gateway devices (requires customized routing, VPC peering) in order to observe and

manage traffic traversing the cloud environment, and they can have multiple ENIs to tap into multiple subnets.

Traffic Mirroring

Traffic mirroring in the cloud used to be challenging, requiring the installation and management of third-party agents on Amazon EC2 instances to capture and mirror EC2 instance traffic. One such platform is Gigamon's GigaVUE CloudSuite for AWS, which acquires, optimizes and distributes selected traffic to security and monitoring tools by performing traffic acquisition using G-vTAP agents.

Amazon VPC Traffic Mirroring addresses these challenges and enables customers to natively replicate their network traffic without having to install and run packet-forwarding agents on Amazon EC2 instances. Amazon VPC Traffic Mirroring captures packets at the ENI level, which cannot be tampered with from the user space, thus offering better security. It also supports traffic filtering and packet truncation, allowing selective monitoring of network traffic. AWS Marketplace already has monitoring solutions integrated with Amazon VPC Traffic Mirroring, such as ExtraHop Reveal(x) Cloud.

The main elements of VPC traffic mirroring are:

- **Mirror source** — An AWS network resource (ENI) in a VPC
- **Mirror target** — An ENI or network load balancer that is the destination for the mirrored traffic
- **Mirror filter** — A set of rules that defines the traffic that is copied in a traffic mirror session
- **Mirror session** — An entity that describes traffic mirroring from a source to a target using filters

The mirror target can be in the same AWS account as the mirror source or in a cross-account AWS environment, capturing traffic from VPCs spread across many AWS accounts and then routing it to a central VPC for inspection. The filter can specify protocol, source and destination port ranges, and classless inter-domain routing (CIDR) blocks for the source and destination. Rules are numbered and processed in order within the scope of a particular mirror session. Sessions are also numbered and evaluated in order. The first match (accept or reject) determines the fate of the packet, because a given packet is sent to at most one target.

Be aware that VPC traffic mirroring is unlike a traditional network tap or mirror port. Mirrored traffic is encapsulated with a VXLAN header and then routed by using the VPC route table. VXLAN traffic (UDP port 4789) must be allowed from the traffic mirror source in the security groups that are associated with the traffic mirror target. Applications that receive the mirrored traffic should be able to parse these VXLAN-encapsulated packets.

Amazon VPC Traffic Mirroring is a game-changer that opens up the possibility of bringing traditional network security monitoring solutions into the AWS environment.

Host-Based IDS/IPS

On the other side of IDS/IPS are host-based IDS/IPS (HIDS/HIPS) and anti-malware solutions. The good news is that these tools can be installed on cloud virtual machines in the same way as on premises. Note, however, that most traditional HIDS/ HIPS agents require more resources, which usually comes with a performance impact on the instances.

Host security monitoring also tends to be more complex to manage. Sensors/agents must be deployed so that they can report back to a management server for analysis. Security teams must take care of event management and log collection and consider network bandwidth to decide whether they want to send the events back to on-premises systems, another virtual machine instance in AWS or maybe to another (SaaS) cloud service. Every time a new instance gets brought up or terminated, the security team must make sure the sensor/agent has to be deployed or decommissioned properly.

Fortunately, there are more cloud-focused, integrated HIDS/HIPS and anti-malware marketplace offerings, such as Trend Micro Deep Security, CloudPassage and Dome9 (now part of Check Point), that can be distributed at the hypervisor layer. Next-generation antivirus (NGAV) and EDR tools like Carbon Black or CrowdStrike have also moved to a SaaS model to support cloud deployments.

Event Management and Analysis

After identifying the most important data sources, collecting data from them and deploying security sensors, we need the means to manage the data collected. Event management and monitoring in a cloud environment consist of activities like scanning for vulnerabilities, event monitoring, alerting, correlation and analysis.

Many security analysts are aware of Amazon CloudWatch, a monitoring and management service available within AWS. Amazon CloudWatch is a highly flexible, general-purpose tool that is not only

meant for security, but also to get a unified view of operational health by monitor applications, resource utilization or systemwide performance changes.

Amazon CloudWatch basically functions as a repository for logs and metrics. AWS services put metrics into the repository, and statistics can be calculated based on those metrics. This statistical data can then be displayed graphically with visualizations (graphs) and dashboards. There are many default metrics available, and custom metrics can be defined too.

Amazon CloudWatch can take logs from Amazon EC2 instances (CPU, memory, network usage, etc.) every five minutes (basic monitoring) or every minute (detailed monitoring), and it has agents that can be installed on instances to send their operating system logs. Amazon CloudWatch Logs can also be used to store and analyze logs from AWS CloudTrail and Amazon VPC Flow Logs. These log entries can be filtered into metrics to define alarms.

The most significant benefit of Amazon CloudWatch is that it is very well integrated with almost every other AWS service, including AWS Security Hub. You can create alarms and periodic events and send them to other tools (for example, AWS Lambda or Amazon Simple Notification Service [Amazon SNS]), and make automatic changes to the resources you are monitoring when a threshold is reached.

AWS Security Hub consumes data from services like AWS Config, Amazon GuardDuty, Amazon Inspector and Amazon Macie, and from supported APN Partner Solutions. AWS Security Hub reduces the effort of collecting all this information. It provides a single, comprehensive view that aggregates, organizes and prioritizes security alerts using a consistent findings format. These findings are displayed on dashboards with actionable graphs and tables.

Putting It All Together

AWS offers various services and access to security, identity and compliance tools from AWS partners. These include firewalls, network or endpoint IDS/IPS applications, and vulnerability and compliance scanners. Because they can easily generate thousands of security events and alerts every day, all in different formats and stored across different platforms, a unified interface is needed for management. Figure 3 illustrates what that unified interface should include.

Amazon GuardDuty is an AWS threat detection service that continuously monitors for malicious activity and unauthorized behavior. The analysis is based on threat intelligence feeds (such as lists of malicious IPs, domains, URLs from Amazon GuardDuty partners) and machine learning to identify unexpected, potentially unauthorized and malicious activity.



Figure 3. Unified Interface for Management of Events and Alerts

Amazon GuardDuty combines, analyzes and processes the following data sources:

- **AWS CloudTrail event logs** — Monitors all access and behavior of AWS accounts and infrastructure
- **Amazon VPC Flow Logs and DNS logs** — Identifies malicious, unauthorized or unexpected behavior in AWS accounts and infrastructure

It is important to note that Amazon GuardDuty was not designed to manage logs or make them accessible in your account. It is built for AWS workloads and AWS data, and is not intended to support data from on-premises or other services. For example, in the case of DNS logs, Amazon GuardDuty can access and process DNS logs through the internal AWS DNS resolvers, but not from third-party DNS resolvers. After it receives the logs, it extracts various fields from these logs for profiling and anomaly detection, and then discards the logs. It is important to collect and store your flow and API logs, as discussed in the “Data Collection” section, in order to retain them for investigations.

The produced security findings (potential security issues) can be viewed in the console, retrieved via an HTTPS API. Alternatively, Amazon GuardDuty can create Amazon CloudWatch Events that can be sent to a SIEM platform, or automated remediation actions can be performed by using AWS Lambda.

Security findings are assigned a severity level of high, medium, or low. These findings are detailed and include information about the affected resource as well as attacker IP address, ASN and IP address geolocation. Amazon GuardDuty has various finding types that cover the entire attacker life cycle, such as reconnaissance, unauthorized access, privilege escalation and persistence.

By importing these findings into AWS Security Hub, you can filter and archive results and create a collection of findings, called “insights,” that are grouped. Insights help to identify common security issues that may require remediation action. AWS Security Hub includes several managed insights by default, but you can create custom insights too. These findings are displayed on dashboards with actionable graphs and tables.

AWS Security Hub also generates its own findings by running automated, continuous configuration and compliance checks based on industry standards and best practices from the Center for Internet Security (CIS) AWS Foundations Benchmark, which is enabled by default. These checks provide a compliance score and identify specific accounts or resources that require attention.

To take advantage of the benefits AWS Security Hub provides, you have to enable and configure the settings of these data sources through their respective consoles or APIs. AWS Security Hub also integrates with AWS CloudTrail, which captures API calls for AWS Security Hub as events.

Organizations may need to use additional third-party tools to integrate with existing tools, to meet compliance requirements or simply to leverage additional features. AWS partners have several cloud-focused event management platforms available. Sumo Logic is a cloud-native data analytics platform, not only for security, but also for operations and business usage. Sumo Logic offers SIEM functionality and machine learning analytics to create baselines and perform anomaly-based detection. Splunk Technology also has several tools for cloud event management, such as Splunk Cloud for security and operational visibility. Open source analytics and monitoring hosted offerings like Amazon Elasticsearch Service on Elastic Cloud and Grafana are also available in AWS Marketplace. Alternatively, Amazon Elasticsearch Service offers Elasticsearch, managed Kibana and integrations with Logstash and other AWS Services.

Automation

The final step in the threat detection strategy is to bring in tools to automate response and remediation after the detection of a threat or vulnerability. This model has three major components:

- **Collecting and monitoring for events** occurring in the environment using AWS CloudTrail logs, Amazon VPC Flow Logs and Amazon VPC Traffic Mirroring. Automated assessment services such as Amazon Inspector, CloudPassage Halo or AWS Config can collect security audit results.
- **Triggering alerts** based on specific patterns and anomalies by relying on Amazon CloudWatch alarms, Amazon GuardDuty findings or alerts from third-party SIEMs. Amazon SNS can be used together with Amazon CloudWatch to send messages when an alarm threshold is reached.

- **Taking action** and starting an automated reaction with tools like AWS Lambda. AWS services such as Amazon CloudWatch or Amazon GuardDuty can automatically trigger AWS Lambda code to perform actions. AWS Systems Manager also has the capability to run automation workflows with triggers using AWS Systems Manager State Manager. Security teams can also take advantage of security orchestration, automation and response (SOAR) platforms like Splunk Phantom or Palo Alto Demisto.

Now, in the next section, we bring together all the steps in building a threat detection strategy.

Security Monitoring Best Practices in AWS

A security team that takes into consideration the recommendations of the previous sections and makes the time investment to fit together the different detection components is able to use cloud-native services and define automated detection and remediation workflows. By reducing the amount of manual labor in the team, the team has more time to focus on other areas of information security.

“By reducing the amount of manual labor in the team, the team has more time to focus on other areas of information security.”

AWS Security Monitoring Best Practices

Some of the most important security monitoring recommendations for the team include:

- Turn on AWS CloudTrail logging in every Region and integrate it with Amazon CloudWatch Logs. Ensure that log file validation is enabled and that logs are encrypted using AWS Key Management Service (KMS).
- Turn on Amazon VPC Flow Logs for every VPC, or at least for the ones with critical assets.

- Leverage Amazon S3 bucket versioning for secure retention and use Object Lock to block object version deletion. Create Write-Once-Read-Many Archive Storage with Amazon S3 Glacier for long-term storage.
- Aggregate AWS CloudTrail log files from multiple accounts to a single bucket. It is a good security practice to set up a separate account and replicate logs to that account, so logs cannot be deleted for a particular account.
- Monitor events and set up Amazon CloudWatch alarms for:
 - User and identity and access management (IAM) activity, especially login events and admin user activity
 - Resource creation events
 - Failed access to resources
 - Policy and configuration changes
 - VPC configuration changes related to security groups, NACs, network gateways, route tables, etc.
 - API calls such as storage attribute changes, unauthorized calls and AWS Lambda events
 - Activity in unusual Regions and at unusual time frames

The CIS has benchmarks on AWS monitoring and logging, offering basic but sound recommendations that anyone can implement and use as a starting point:

- The **CIS Amazon Web Services Foundations** document provides guidance for configuring security options for a subset of AWS.
- **CIS Amazon Web Services Three-tier Web** provides guidance for establishing a secure operational posture for a three-tier web architecture deployed to the AWS environment.

The Process

This process has to start with data collection. The security team must set up AWS API and user activity logging with AWS CloudTrail. These logs are the team's sources for the metrics and triggers it identifies for the Amazon CloudWatch alarms. This already makes the team capable of responding automatically to events such as resource changes, for example, when someone tries to disable AWS CloudTrail logging or log in with an AWS account root user at unexpected times from an unexpected location. These can be simple rules to indicate the events of interest and the automated actions to take when an event matches a rule. The actions that can be triggered include but are not limited to:

- Invoking an AWS Lambda function
- Invoking Amazon EC2 Run Command
- Notifying an Amazon SNS topic

To monitor network traffic and packet flows in its VPCs, the team will rely on Amazon VPC Flow Logs and configure Amazon VPC Traffic Mirroring to send traffic from instances to network security sensors. Depending on the skill set of the security team members, the team might choose to use open source tools for its NIDS/NIPS and HIDS/HIPS needs, or deploy APN partner AMIs like NGFW/UTM appliances across their VPCs.

If the security team wants to go one step further, it can enable AWS-built services like AWS Trusted Advisor, AWS Config, Amazon Inspector and Amazon GuardDuty. These are designed to exchange data and interact with other core AWS services, to identify potential security findings and raise alarms.

AWS Security Hub or an APN partner event management service could allow the team to enable, configure and connect APN partner tools and review findings in one central location. AWS Security Hub can also automatically send all findings to Amazon CloudWatch Events. After an Amazon CloudWatch Event is sent or a finding notification is posted to an SNS topic, an AWS Lambda function can be triggered, and services like AWS Systems Manager can be used from within the AWS Lambda function to perform automatic remediation on the instance.

Conclusion

By relying on the most common data sources, organizations can build a powerful threat detection strategy and gradually improve their monitoring capabilities. The focus should be on the data types that can provide the highest value and not only cover network and system monitoring but also have the information needed for cloud environment monitoring. Advancements in monitoring, such as Amazon VPC Traffic Mirroring, can be the means of adapting traditional security monitoring techniques to the cloud.

Collecting the data is just one half of the equation. Without analysis and event management, data collection does not provide any value. Analysts can detect suspicious or malicious events during a manual threat hunting process or alerts could be triggered based on predefined conditions, rules or machine learning. Combining the different cloud-native services and features available can also help in detecting threats.

The ultimate goal is to take advantage of automation tools that can serve as a force multiplier and assist security teams immensely in incident response and vulnerability remediation by automating the most common tasks.

About the Author

David Szili is a SANS instructor for SANS FOR572: Advanced Network Forensics: Threat Hunting, Analysis, and Incident Response. A managing partner and CTO at a Luxembourg-based consulting company, he has more than eight years of professional experience in penetration testing, red teaming, vulnerability assessment, vulnerability management, security monitoring, security architecture design, incident response, digital forensics and software development. David holds several IT security certifications, including the GSEC, GCFE, GCED, GCIA, GCIH, GMON, GNFA, GYPC, GMOB, OSCP, OSWP and CEH. He is also a member of the BSides Luxembourg conference organizing team.

Chapter 17: How to Perform a Security Investigation in AWS



Kyle Dickinson

SANS Instructor & Author

"One of the more common questions that I receive is, 'Now that we've moved to the cloud, how does the investigation or incident response process change?'"

Throughout this chapter I address the different cloud-specific considerations teams need to review throughout the SANS incident response steps, and what they can do to improve and update their processes when moving to the cloud—because as they will discover, not all of the processes that they have used traditionally for on-premises incident response and investigations necessarily migrate to the cloud."

Introduction

With the rapid growth of cloud service providers and the appeal, for organizations, of no longer having to manage their own data centers, more organizations are migrating to infrastructure-as-a-service (IaaS) providers. And the ability to stand up global infrastructure in a few clicks, or through a Continuous Integration and Continuous Deployment (CI/CD) pipeline, is drawing developers to cloud services as well.

What does this mean for incident response and forensics teams? We advocate for putting cloud-specific plans into place, because the technologies that enable investigations in the cloud differ from the ones for on premises, as do the levels of responsibility.

In this paper, we cover incident response plans in IaaS implementations, various services available that aid in conducting an investigation and the different components of an audit log. We also explore how to perform a forensic image analysis and how to review the communications that are coming to and from an EC2 instance.

Investigations vs. Incident Response

Investigations (or forensics), by definition is "... the process of using scientific knowledge for collecting, analyzing, and presenting evidence. ..." ¹ Although investigations do not have to be aimed at providing evidence for a court case, understanding the process is important. We examine these two data sources next.

Investigations

The process of using scientific knowledge to collect, analyze and present evidence

Incident response

The process of using knowledge gained from an investigation to address a security incident

¹ US-Cert, "Computer Forensics," www.us-cert.gov/sites/default/files/publications/forensics.pdf

How Investigations Differ in Cloud-Based Environments

When performing an investigation in Amazon Web Services (AWS),² it's essential to understand that the investigation "playbook," or process, that an organization has for on-premises investigations is not exactly the same as for cloud-based investigations. Table 1 shows the differences between on-premises and cloud-based investigations.

The majority of the data sources and preparatory steps should be included in an incident response plan, which changes based on the type of cloud service model that is being consumed, such as software-as-a-service (SaaS), platform-as-a-service (PaaS) and infrastructure-as-a-service (IaaS).

Table 1. On-Premises vs. Cloud-Based Investigations		
Process	On-Premises	In the Cloud
Disk imaging	Physical drive connected to forensic workstation	Snapshot taken from Amazon EC2 instance, converted to volume and attached to forensic instance
Memory acquisition	Physical access to workstation as it's running	Private key or local user/trusted host access required
Network logging	PCAP in-line with netflow	Amazon VPC Traffic Mirroring

The Incident Response Process

Let's start by outlining the incident response process. An incident response is typically triggered by reports of "something happening" or notification that "something happened." Figure 1 shows the step for responding using the SANS six-step incident response methodology.³ This methodology can easily be adapted to cloud-based environments.



Figure 1. SANS Incident Response Steps

² Because this paper is an exploration of performing investigations in AWS, it is important to talk about the tools available. The use of these examples is not an endorsement of any product or service.

³ "Incident Handler's Handbook," December 2011, www.sans.org/reading-room/whitepapers/incident/incident-handlers-handbook-33901

Here's a simple example:

Preparation

- What cloud service provider is being used?
- What is the deployment model? (Public, hybrid, private?)
- What is the cloud model? (SaaS, PaaS, IaaS?)

Identification

- Is there unusual activity in the audit logs?
- Did something get misconfigured?

Containment

- Can we disable a user's access?
- Can we isolate the VM or subnet?
- How do we acquire an image?

Eradication

- Can we remove affected systems?
- Can we remove/replace compromised credentials?

Recovery

- Can we restore normal business operations?
- Is a business continuity plan available?
- Did that plan need to be implemented?

Lessons Learned

- What gaps in coverage did we discover?
- How do we close those gaps?

For cloud-based environments, the preceding methodology does not provide a complete incident response plan; however, we can see there may be some crossover from an on-premises plan, but it is not a one-for-one replacement when moving to the cloud.

Shared Responsibility Model

The shared responsibility model is a common method of determining where the responsibility shifts and which party is responsible for specific parts of the infrastructure. Depending on the type of service you’re consuming, the provider can be responsible for some aspects or most aspects of the cloud.

Typically, with IaaS, the provider is responsible for security of the cloud, while our security teams are responsible for security in the cloud. When moving to IaaS providers, such as AWS, security teams must consider capabilities and services like the ones shown in Table 2.

Table 2. Key Capabilities and Services		
Capability	AWS Service	Description
Compute	Amazon Elastic Cloud Compute (EC2)	Uses Amazon Machine Images (AMIs) to get started Multiple OS support Pay for what you use Next-gen Nitro infrastructure, created by AWS
Storage	Amazon Elastic Block Store (EBS), Amazon Simple Storage Service (S3), Amazon Elastic File System (EFS)	Amazon S3 offers multiple storage classes for multiple use cases. Amazon EBS is used for the “block device” or hard drive for Amazon EC2 instances. Amazon EFS is used for file sharing storage with two storage classes to choose from.
NetFlow	Amazon VPC Flow Logs, Amazon VPC Traffic Mirroring	Capture information of network traffic going in and out of a VPC
Auditing	AWS CloudTrail	User attribution data Log integrity can be enabled Can send data to an Amazon S3 bucket for storage/archival

Modern Security Controls

A typical on-premises environment may include the following tools that could be used in conducting incident response or investigations:

- Network intrusion detection systems (NIDS)
- Packet capture devices or network taps
- Vulnerability management scanners
- Endpoint detection
- Proxies and firewalls

When we move our investigations to a cloud-based environment, there are no decisions like “Where to ship my NIDS, network taps, vulnerability management, etc. ...” details. This is because we lose physical access to our infrastructure. That is okay. Instead of worrying about physical infrastructure, we can now focus on how to modernize our security controls.

AWS Marketplace allows security teams to stand up modern tooling that can come in the form of SaaS or AMIs and allow organizations to use the capabilities provided by AWS Partners to supplement the services that are available directly from AWS.

To better understand how to conduct an investigation within AWS, it is best that we understand the native services available to security practitioners so that we can understand what is and is not possible out of the box. This also strengthens the understanding of how to integrate the different capabilities that third-party tools offer.

Using AWS Services in Investigations

As part of the evidence gathering and analysis process, user attribution information tells us about the activity that a particular resource or user has performed. In the following sections, we discuss these activities as well as describe how to gain insight into network traffic.

Understanding User Activity

AWS CloudTrail gives security teams the who/what/when/where/how of the activity being investigated. This is the information that the auditing data teams need to better understand a user's actions. By default, AWS CloudTrail is enabled within the AWS Management Console. However, to ship these logs out of the account to a SIEM or log analysis tool, we need to set up a trail first.

If we look at an example of an AWS CloudTrail log in the AWS Management Console, security teams have multiple ways to search for data:

- **Username** — Search by the user's name
- **Event name** — Search by a specific API call (e.g., DeleteTrail)
- **Resource type** — Search by an AWS service type (e.g., Amazon EC2 instance)
- **Resource name** — Search by a resource name (e.g., instance ID, ENI)
- **Event source** — Search results from specific AWS services
- **Event ID** — Search based on a unique ID for an AWS CloudTrail event
- **AWS access key** — Search by access key to show what was done in a single session

Figure 2 shows an example of an AWS CloudTrail event.

By looking at the single AWS CloudTrail event shown in Figure 2, we can piece together that the user (Marc the intern) successfully logged into the AWS Management Console using Google Chrome, from IP address 11.22.33.44, using a password with no multifactor authentication.

Keeping this information in mind, the majority of these fields remain persistent in each AWS CloudTrail event as we look to conduct an investigation. Having this data visualized and stored in a central location aids us significantly. Not only do we benefit from having the logfiles stored in a single location under the security team's control, but we have heightened security controls around this storage. Visualization allows investigators to demonstrate the activity and the location from which the activity was performed.




Figure 2. An AWS CloudTrail Event

“We highly recommend that you enable Amazon VPC Flow Logs for your VPCs; they are not enabled by default.”

Gaining Visibility into Network Traffic

Amazon VPC Flow Logs provide visibility of network traffic going in and out of a VPC, also known as north-south traffic.

Looking at the structure of a VPC Flow Log, we see the details listed in Figure 3.

Amazon VPC Flow Logs give us a high-level view of network traffic. Exporting this data to a SIEM can add more context to Flow Logs by correlating threat intelligence data to the source or destination IP addresses to determine whether Amazon EC2 instances are communicating to potentially hostile hosts, such as those known from cryptomining or botnets.

Amazon VPC Traffic Mirroring is another method of obtaining insight into your network traffic that is available on AWS Nitro instances. What's handy about Amazon VPC Traffic Mirroring is that it's a "spanport-as-a-service" that enables security to send all north-south traffic to another instance for further analysis, if required, or integrate to another traffic-analysis toolset.

Account ID	eni-0fe570007a111e2e3	Source IP	104.248.185.25	Source Port	10.0.1.19	Destination IP	32767	Destination Port	3389	Protocol	6	Bytes Transferred	1	40	Start Time	1567551544	End Time	1567551586	Action Taken	REJECT OK
------------	-----------------------	-----------	----------------	-------------	-----------	----------------	-------	------------------	------	----------	---	-------------------	---	----	------------	------------	----------	------------	--------------	-----------

Figure 3. Structure of a VPC Flow Log

Forensic Acquisition

Should the incident require the security team to perform forensics on an Amazon EC2 instance, we need to take a snapshot of that instance and create a volume from that snapshot to share to a SIFT Forensic Workstation.

The following steps are an example of that process for a compromised implementation:

1. Create a security group that does not allow outbound traffic
2. Attach to compromised Amazon EC2 instance
3. Take snapshot of Amazon EC2 instance
4. Perform memory acquisition, if possible
5. Share snapshot with Security Account (if using one)
6. Create volume from snapshot

7. Attach volume to SIFT EC2 instance

8. Conduct forensics

It is possible to automate this process, which would provide faster data acquisition and response.

About the Author

Kyle Dickinson teaches SANS SEC545: Cloud Security Architecture and Operations and has contributed to the creation of other SANS courses. He is a cloud security architect for one of the largest privately held companies in the United States. As a strategic consultant in his organization, Kyle partners with businesses in various industries to better understand security and risks associated with cloud services. He has held many roles in IT, ranging from systems administration to network engineering and from endpoint architecture to incident response and forensic analysis. Kyle enjoys sharing information from his experiences of successes and failures.

Chapter 18: How to Leverage Endpoint Detection and Response (EDR) in AWS Investigations



Justin Henderson

SANS Certified Instructor & Author

"Detection is a game that is difficult to play well. Visibility and understanding go hand-in-hand when it comes to decision making. Is something suspicious? Why? If you do not have visibility, then your informed decisions are null and void. Endpoint detection and response (EDR) provides in-depth coverage of what is occurring within an operating system and works for private organizations as well as large-scale cloud solutions.

Cloud deployments of EDR solutions support auto-deployment, asset control and dynamic rulesets. Most importantly, they provide context-driven detection that educates and informs organizations. The result is an ability to take control of your assets enterprise-wide."

Introduction

The security challenges organizations face are often a direct result of evolving technologies such as virtual machines, containers, storage and even serverless code. Technology is not static. It changes dynamically via new developments such as infrastructure as code (IaC) and auto-scaling capabilities found at multiple layers of service. The result of this technological evolution is complexity in cloud environments. To secure such environments, you have to know and understand them.

Effective security teams implement appropriate technologies to mitigate potential challenges—for example, EC2 instances configured in a way that allows fileless malware such as the PowerShell Invoke-Mimikatz to steal credentials, or unsecured containers that an attacker can inject a PHP or .NET web shell into in order to access files and databases in Amazon S3 buckets, MySQL or an Amazon Relational Database Service (RDS). To enable more effective approaches to ensuring security, this paper illustrates how to leverage endpoint detection and response (EDR) in Amazon Web Services (AWS) to achieve a higher standard of security while simplifying management overhead. The goal is to ease the burden of cloud security via EDR technologies.

Acquiring Cloud Visibility

The first step in securing an AWS environment is not unique: Security teams need to understand what assets they have. After all, you cannot protect what you do not know exists. Traditionally this is a three-step process, as defined in Table 1.

Table 1. Asset Identification Process

Step	Definition	Example
Network scanning	A process to identify your assets and where they exist	Performing a port scan of Amazon EC2 instances
Service enumeration	A process to identify assets by querying a management service	Asking Kubernetes or Docker what containers exist
Agent installation	A process to push a security agent to an asset	Installing or using a log agent like Syslog-NG

But when it comes to cloud visibility, that traditional approach could leave gaps in coverage because of the way customers configure their environment. Good security practices involve customers locking down their assets, but a network scan would not identify all EC2 instances, because of customer configuration of Amazon security policies, network firewalls, and potentially endpoint controls or configurations. The lockdown of EC2 assets improves security, but it also makes 100% asset discovery difficult or impossible. Yes, an agent can easily be deployed to EC2 instances. However, because of an inability to see all instances and understand the underlying operating system, it is not possible to be aware of all assets in order to push agents to them. A more comprehensive approach is needed.

In addition, containers, Amazon S3 storage and serverless code execution are not traditional computer technologies. For them, deploying an agent is not necessarily an option, and even if it is for your organization, we recommend against this practice. Consider an Amazon EKS container running Nginx. This container is designed to run Nginx and nothing else, as indicated by the following code:

	PID	%CPU	%MEM	VSZ	RSS	TTY	STAT	START	TIME	COMMAND
root	1	0.0	0.1	10632	5488	pts/0	Ss+	18:49	0:00	nginx: master process nginx -g daemon off
nginx	6	0.0	0.0	11104	2664	pts/0	S+	18:49	0:00	nginx: worker process

Can you deploy an agent within a container? Yes. Should you? No, because deploying agents to a container introduces software dependencies, increases computational resources and adds management overhead.

However, without the ability to discover and protect containers, you are exposing yourself to a lot of risks. The same holds true for other services such as Amazon S3 storage. You cannot directly deploy an agent to an S3 bucket, but it still needs to be monitored for unauthorized access.

To achieve a holistic view of your AWS environment, consider adopting a modern methodology that integrates with AWS. AWS supports multiple EDR vendors that utilize Amazon APIs to move past the “everything requires an agent” approach. The steps outlined in Figure 1 on the next page show a more modern process.

Adopting a unified and holistic view of assets brings a simplified understanding of your environment. You can easily deploy these solutions, requiring you only to choose and subscribe to the vendor in AWS Marketplace. For example, subscribing to CrowdStrike’s EDR¹ provides the capability to probe Amazon EC2, Amazon Elastic Container Service (ECS), and Amazon Elastic Kubernetes Service (EKS) to provide EDR, next-generation antivirus, threat intelligence and more.




Figure 1. Modern Asset Identification Process

Deploying Controls to EC2 Instances

When implementing security controls to EC2 instances, it is imperative to plan for scale. What happens when you add or remove EC2 instances?

A good place to begin is with the Center for Internet Security's (CIS) Critical Security Controls¹ 1 and 2: Keep an inventory of authorized and unauthorized hardware and software. An effective AWS EDR strategy incorporates this principle by supporting automatic deployments.

Let's use CrowdStrike's EDR solution² to demonstrate how to integrate EDR in AWS. The process for deploying EDR in AWS using CrowdStrike follows these steps:

1. **Subscribe to CrowdStrike EDR (found in AWS Marketplace).**

2. **Deploy CrowdStrike Falcon Discover.**

- a. Falcon Discover acquires access keys to query AWS. With these keys, it identifies all EC2 instances, even across regions.

¹CrowdStrike, CrowdStrike Falcon and Falcon Discover are trademarks or registered trademarks of CrowdStrike Inc.

²www.cisecurity.org/controls

³This paper mentions solutions to provide a real-life example of how to integrate EDR in AWS. The use of these examples is not an endorsement of any solution.

- b. The user authorizes Falcon Discover to deploy agents to specific EC2 instances or all instances automatically.
- c. Agents continuously auto-deploy to authorized instances.
- d. Optional: Falcon Discover is configured to monitor other assets such as CloudTrail. If enabled, this capability provides additional security controls such as alerting on tenant-level security controls.

3. The organization reports on asset coverage and monitors alerts.

“An EDR solution should auto-scale and grow with you, not slow you down.”

Achieving Proper Security Controls

The phrase “Here be dragons” designates unexplored and potentially dangerous areas. For security professionals, there certainly are metaphorical dragons in EDR and caution is necessary. There are many products that claim to be EDR solutions. Although each of them provides endpoint controls, their depth of coverage and capabilities vary, resulting in different levels of protection.

Let’s explore capabilities a successful EDR solution should provide by considering a plausible attack against an EC2 instance.

Attack Scenario:

Consider this scenario:

An organization is running a Windows EC2 instance with MSSQL services. An attacker is trying to identify critical assets but so far has only a standard account on a different EC2 instance. To escalate

privileges, the adversary runs `setspn` to identify accounts vulnerable to what is commonly referred to as a Kerberoasting. Because MSSQL servers use service principal names (SPNs), the adversary finds the EC2 MSSQL service, pulls down a Kerberos ticket and then uses a password cracker to identify the MSSQL service account password. This account is then utilized to gain access to the EC2 instance using `psexec`. From there, the attacker establishes persistence by creating a digitally signed Microsoft executable due to a flaw from missing the patch for CVE-2020-0601, which allows abuse of the cryptographic process for Elliptic Curve Cryptography (ECC) handled by the Windows operating system. That process results in a persistent command and control that looks normal because the binary has been digitally signed by Microsoft. The further activity includes enumerating the MSSQL database.

The scenario provided is a bit convoluted. However, each step utilizes known attack techniques classified by the MITRE ATT&CK framework.⁴ But just because something is a known technique does not mean it automatically should be blocked or flagged as an automatic alert. Consider the breakdown of this scenario:

MITRE T1208 Kerberoasting—`setspn`, `klist` and PowerShell can be utilized to export a Kerberos token. This can then be password-cracked if the password is weak.

- **Identification**—Commands like **`setspn`** are not utilized by standard users and would often be an anomaly.
- **Problem**—System administrators do use **`setspn`**. Alerting on each use would generate multiple false positives.

MITRE T1035 and T1050—The use of **`psexec`** to gain remote access would trigger a new service and its corresponding execution.

- **Identification**—**`psexec`** is not necessary if organizations use other remote access tools, such as PowerShell remoting.
- **Problem**—Organizations may utilize **`psexec`** as a standard remote access tool.

MITRE T1116—Abusing CVE-2020-0601 to create a binary that appears to be digitally signed by Microsoft and then using that binary for persistent callbacks provides an adversary stealth communication.

⁴<https://attack.mitre.org>; MITRE ATT&CK Matrix is a trademark of The MITRE Corp.

- **Identification**—A digitally signed certificate should conform to proper Elliptic Curve Cryptography (ECC) standards.
- **Problem**—Software may use different algorithms, key lengths and other attributes when generating certificates.

MITRE T1219—The adversary left a binary on the MSSQL server to maintain remote access.

- **Identification**—Persistence mechanisms generate network traffic to other assets that should not be happening.
- **Problem**—Because of asset management, patching and other system processes, it may be difficult to distinguish a good network callback from a malicious one.

Given this scenario, a proper EDR solution should provide multiple angles to identify the adversary. Each step could be a regular event. However, by analyzing the series of events, an EDR solution should clearly identify and even stop this attack. The following sections describe the features to look for in modern EDR solutions that would aid in this attack.

Process Tree

One method of finding unwanted activity is monitoring each process. This includes process, command line, parent process, parent process command line, user, integrity level and other related variables. This information then is correlated with the chain of processes occurring. EDR should identify abnormal processes or an abnormal chain of events and provide a visual process tree to explain why something is considered harmful (see Figure 2).

“Organizations need to cautiously evaluate EDR solutions against modern threats and risks.”




Figure 2. A Process Tree Diagram

MITRE Tagging

Instead of reinventing the wheel, EDR solutions should integrate with known, proven frameworks. The MITRE ATT&CK framework is one of the most practical approaches to identifying attacker techniques, tools and behaviors. Each piece on its own is not enough to block an attack or generate an alert. However, specific techniques are more likely to be malicious than others, and EDR solutions can search for a combination or sequence of techniques and score them. Commercial EDR scores then combine to block or identify an attack, plus help analysts by telling a story of what happened. Figure 3 provides a sample visualization of the attack.




Figure 3. Graphical Description of the Attack

Signatures, Heuristics and Machine Learning

EDR should include domain expert-based heuristics as well as potential algorithms that adapt over time, such as supervised or unsupervised learning. In the sample scenario, a basic heuristic check would identify that Kerberos reconnaissance commands were issued, followed by an authentication request from the original source EC2 instance. Machine learning may identify that the source user is highly unlikely to run commands like **klist** or **setspn**. Even traditional signatures may work by looking for an improperly formed Elliptic Curve Cryptography (ECC) generator set that abused CVE-2020-0601.

IoC Support

A robust EDR solution should offer the ability to identify a given activity and search for it across the entire environment. Put plainly, an organization should be able to identify the characteristics of an attack and document them in the form of an indicator of compromise (IoC). IoCs can be specific and straightforward, such as the **SHA1** of the binary used for persistence in the scenario. Or they can be specific, with broad characteristics such as looking for certificate files with specific algorithms, key lengths and file sizes.

Ideally, IoC support should include vendor-defined IoCs that regularly update, plus the ability to develop internal IoCs and perform threat hunting with them. With such capabilities, organizations can perform investigations looking for IoCs from previously identified IoCs or proactively by looking for IoCs shared from external parties. Support for standard IoC formats such as **YARA** should be given consideration so that IoCs work outside the EDR platform.

Provide Attribution

Attribution is the ability to associate something with a person or entity. Within EDR, organizations should utilize MITRE ATT&CK and any proprietary sources to help them understand who or what is attacking them. At a minimum, such information is useful to understand what may occur next. For example, with profiling, various techniques, tools and IoCs may indicate that a known threat group is in play. In our scenario, profiling may inform the organization that the specific attack group has access to the EC2 instance, and the organization should look for specific backdoor programs. More importantly, the information can predict what the attack group's goal is, such as stealing healthcare information. Using this, an organization can make an informed decision to pause the EC2 instance or take alternative steps.

Response Capabilities

Given enough high-fidelity information, EDR should block or reverse the damage from the attack. If the attack was ransomware, EDR should restore encrypted files pre-ransomware. In our scenario, given that it is possible that the attack would not be blocked until a certificate was generated to exploit CVE-2020-0601, EDR would identify the attack and notify an analyst. Then, an analyst could choose to take remediation actions given in prior steps in the scenario. The response does not have to be the standard

“EDR should be a sum of its parts: signatures for known bad, heuristics based on domain-expertise and machine learning for finding anomalies.”

one of blocking a connection and removing a file. A response should mean taking steps against the full scenario—for example, removing the persistence file abusing CVE-2020-0601, killing the psexec process and killing the process providing remote access to the initial EC2 instance.

Real-Time Vulnerability Reporting

Because an EDR solution resides directly on an endpoint, it should identify all software that is installed or running. Because of this, it is possible to have real-time vulnerability reporting. Vulnerability scanning is hard to scale, but with an EDR partner that uses it for vulnerability reporting, it does not have to be.

EDR and Container Security

EDR solutions often employ agents for robust operating system visibility and protection mechanisms. But what about other deployments such as Amazon ECS and Amazon EKS containers? Some EDR solutions have no coverage for containers or anything that is not a traditional endpoint.

An EDR deployment in AWS should provide coverage to more than just EC2 instances. Fortunately, multiple vendors support a broader range of coverage in the AWS cloud. Regarding containers, the following foundational constraints need consideration.

- **Containers, ideally, should run a single service.**
 - a. Be sure to design a container around a single process.
 - b. Subprocesses such as an Nginx container running a master and worker are inline with best practices. Running multiple processes in parallel is not.
- **Containers should include only software that is necessary.**
 - a. Adding an agent bloats container images.
 - b. Adding an agent also increases overhead computation, thus increasing costs.

Knowing the principles behind deploying and managing containers, deploying an agent is far from ideal. While technically an agent can be embedded into an image, it's a horrible idea due to the agent breaking the aforementioned foundational constraints. Still, most of the attack scenario described previously can work within containers, so if you use containers, be sure you identify an EDR solution that covers containers.

The implementation of EDR into containers requires software that can see into a container. Similar capabilities such as monitoring processes, files and network connections need to work inside a container. To accomplish this, either an agent needs to be deployed into each container, or an image, a sidecar container or a centralized agent needs to be implemented. Let's explore each option:

- **Agent**—Installing agents inside a container is against good practice, hard to manage and computationally expensive.
- **Sidecar**—A sidecar is a concept of deploying a container next to another container, similar to a motorcycle with a sidecar attached. In this case, the sidecar container receives access to the original container so it can monitor it. Technically this option works, but it adds additional computing resources and overhead to ensure each container gets a sidecar.
- **Centralized agent**—A better approach is to have one or more specialized agents that utilize Amazon APIs and access to dip into containers and corresponding images. For example, CrowdStrike EDR supports deploying a single instance of Falcon Insight. Falcon Insight then acts as a centralized agent that interfaces with Amazon ECS and Amazon EKS to secure containers and images.

Using an EDR solution that supports AWS integration dramatically simplifies deployment and ensures minimal gaps in security controls. A centralized agent such as Falcon Insight would identify the CVE-2020-0601 vulnerability in an MSSQL Server image or notify the analyst that the image is no longer vulnerable, but active containers still are. In addition, containers do not run full operating systems, and an EDR solution can more readily apply heuristics and anomaly detection. For example, an MSSQL container should only be running MSSQL. If a binary began a persistent callback mechanism, an EDR solution should be able to intervene to detect and block it.

While all assets eventually are decommissioned, containers are decommissioned much more so. Their ephemeral nature introduces new challenges that only a modern EDR can solve. Consider an MSSQL container that gets infected but later is stopped due to scheduled maintenance. After the maintenance, a new container is deployed without any known vulnerabilities. The problem is the old container included crucial forensics evidence regardless of the compromise. A reliable EDR solution would provide a way

to access terminated containers in order to provide analysis in an ad hoc or as-needed basis. If data was stolen in the prior scenario, the solution could launch an investigation that analyzes the previously decommissioned container.

EDR Integrations: A Platinum Experience

EDR provides multiple angles of coverage from native AWS integration, asset knowledge, and detection and prevention capabilities, up to threat hunting and intelligence.

Because of the extensive visibility capabilities and IoC support, organizations should consider EDR for a third-party integration. What if a breach occurred and data and/or malware was transferred into an S3 bucket or later shifted to an external SaaS provider, such as Dropbox? With data moved outside the endpoint, EDR protection generally stops. Yet some EDR solutions go the extra mile and support integration with other solutions, such as cloud security providers like Netskope.⁵

Instead of running multiple security solutions in parallel, they can be integrated. Think of this as a platinum experience, going above and beyond. Via AWS Marketplace, organizations can quickly subscribe and deploy multiple solutions. Then via partner sharing and documentation, they can quickly integrate multiple products into Amazon's APIs as well as from partner to partner APIs. The result is a streamlined solution with extended coverage.

As an example, consider the use of CrowdStrike and Netskope integration. The two solutions support integration and sharing of IoCs. They also support dynamic access control lists as a result. An IoC showing malware or files stolen can be shared as an IoC in CrowdStrike to Netskope and help identify where the files were staged or moved within multiple cloud tenants. Or maybe the attack never would have succeeded. In the scenario described earlier, the adversary first had to get onto the initial EC2 instance before pivoting to the MSSQL Server. If the first EC2 server was missing CrowdStrike's EDR agent, then a dynamic access control could limit cloud access via the CrowdStrike and Netskope integration. This control may also limit or identify the attacker trying to access or stage payloads.

⁵Netskope is a trademark of Netskope Inc.

Conclusion

The definition of an endpoint is evolving. Endpoints are moving past EC2 virtual machines, and it is imperative for EDR solutions to evolve and support this evolution. AWS is quickly adopting new methodologies of implementing and deploying endpoints as well as technologies such as infrastructure as code. As a result, organizations must understand the gaps and risks of not knowing and understanding the various endpoints found in their AWS infrastructure. Organizations should consider an EDR solution that provides advanced controls and works with their AWS environment rather than around it.

Organizations should choose an EDR that encompasses the multiple types of endpoints, such as Amazon EC2, Amazon ECS and Amazon EKS. Because of other infrastructures, such as containers, EDR needs to move past the mantra of every asset getting an agent. New methods such as centralized agents with Amazon API integration are required to come close to 100% asset coverage remotely. Asset coverage and security controls are further extended with EDR solutions that integrate with other partners via API hooks.

About the Author

Justin Henderson is a certified SANS instructor who authored the SEC555 SIEM with Tactical Analytics course and co-authored SEC455: SIEM Design and Implementation and SEC530: Defensible Security Architecture and Engineering. He is a passionate security researcher and security consultant with over a decade of experience in consulting and is one of the co-founders of H & A Security Solutions. Justin is the 13th of 20 GSEs to become both a red and blue SANS Cyber Guardian and holds 61 industry certifications. He specializes in threat hunting via SIEM, network security monitoring and ad hoc scripting.

Chapter 19: How to Build a Threat Hunting Capability in AWS



Shaun McCullough

SANS Instructor

"The public cloud can significantly change the approach to threat hunting in your environment. Organizations may find that they no longer have the same level of fidelity of log data that they are used to, but also have new tools to gather insights that may have been difficult in their on-premises environments. This chapter focuses on how to build a threat hunting program that is tailored to public clouds, investigate new ways of collecting data, and use specific AWS tools to analyze, detect and respond to threats."

Introduction

The infrastructure is built, a patching plan is in place, firewalls are locked down and monitored, assets are managed, and the SOC team is responding to alerts from the security sensors. When basic security hygiene is implemented, the threat hunting team needs to start evaluating infrastructure for any risks and undetected unauthorized broad access.

Because infrastructures are complex, with many moving parts, teams need a plan to manage all the data from all the various operating systems, networking tools and custom applications. They also need to know which threats to look for, how to prioritize them and where to start hunting.

Cloud environments bring their own set of complexity and peculiarities for threat hunting. Customers realizing the benefits of elastic environments may find that systems that had a threat on Friday are terminated on Sunday. Reliance on cloud services likely means relying on the data they offer in a platform-specific format.

In addition to the cloud, the management plane is now a new threat vector that teams have to consider, along with web apps, virtual machines and databases.

In this paper, we walk through the threat hunting process and how it should fit into an organization's overall security strategy. We discuss how to determine what data to gather, options for analyzing it and the kinds of tools threat hunters can use in cloud environments.

Threat Hunting:

The proactive evaluation of the infrastructure operations to detect a threat beyond the deployed security tools

Threat Hunting on Premises vs. in the Cloud

It is vital to understand the process of threat hunting and how to approach it differently than standard security operations. Let's look at this process in the context of a web application. To enhance understanding, this paper references a common use case found in cloud architecture: managing a web application.

Web Application Use Case

A database-based web application is running and is internet-facing. The virtual machine (VM) is running a critical business application and would be considered a potential target. Although the methods of attack against web applications in the cloud are similar to those on premises, threat hunters must adjust their approach and adopt a new set of tools for detection and remediation.

The cloud management plane is an attack vector that threat hunters must evaluate. If attackers were to gain a foothold in a web application, could they leverage it to get further into the cloud infrastructure? Could they make changes, set up persistence and spin up a cryptocurrency mining rig that will run at great expense to the affected user?

The damage can be financially and legally impactful. The web application is running on an Amazon Elastic Compute Cloud (EC2),¹ a VM, that reaches out to an Amazon S3 bucket to retrieve configuration files every time the server starts up. This use case, illustrated in Figure 1, is simplified by design to help tell the threat hunting story. A properly architected web application would include additional protections..




Figure 1. Web Application Use Case

¹ This paper mentions product names to provide real-life examples of how threat hunting tools can be used. The use of these examples is not an endorsement of any product.

How to Approach Threat Hunting

Threat hunting is more of an art than a science, in that its approach and implementation can differ substantially among various organizations and still be right. Every organization builds and operates its infrastructure in its own way; their teams have varied compositions of skill sets, talents and goals, and they face different threat risks.

CIS Critical Controls Are Vital to Threat Hunting	
The Center for Internet Security (CIS) identifies 20 essential security controls, the first six of which are basic controls. Table 1 lists these basics controls and describes their importance to creating an effective threat hunting program.	
Table 1. CIS Critical Controls and Threat Hunting ²	
CIS Control	Description
Control 1: Inventory and Control of Hardware Assets	Threat hunters need to know and manage hardware and software assets, so they can identify which infrastructure services to evaluate and what software is approved.
Control 2: Inventory and Control of Software Assets	
Control 3: Continuous Vulnerability Management	By eliminating software vulnerabilities, threat hunters can save time and resources.
Control 5: Secure Configuration for Hardware and Software on Mobile Devices, Laptops, Workstations and Servers	
Control 4: Controlled Use of Administrative Privileges	Organizations should limit the use of admin privileges so threat hunters can better determine what is legitimate use.
Control 6: Maintenance, Monitoring and Analysis of Audit Logs	The core of threat hunting relies on proper managing, monitoring and analysis of logs.

Threat hunting is about approaching security from a different angle. For instance, the security operations center (SOC) has a collection of alerts from various security products, such as antivirus scans, email security solutions, vulnerability scans, firewall alerts, IDS/IPS, and login failures. If a scan shows that a production server is vulnerable with a critical alert, a SOC member creates a ticket for the serveradministration teams to plan for an update. The driver of that interaction is a security product alerting on a strong indicator. Thus a workload needs to be patched.

² www.cisecurity.org/controls/cis-controls-list/

Threat hunting starts with the premise of, “Our main web application is facing the internet and may be the victim of a web attack. Let’s see how we can determine that.” Or maybe a weak indicator sparks suspicion:

“Multiple failed SQL injection attacks in a row. The web server performance is slower. Let’s look for potential intrusions.”

There are multiple scenarios in between that can all be considered threat hunting. With a strong indicator from a security service, there is a process in place to remedy the situation. With threat hunting, the team is looking for anomalous behaviors without strong indicators. The outcome is likely unknown, the investigation is murky, and the process is research intensive. It is essential to build a threat hunting process and environment to maximize the effectiveness of the team.

Threat Hunting Loop

Building a threat hunting process from scratch takes time, resources and the ability to reach out to experts inside and outside the organization. The Threat Hunting Loop,³ shown in Figure 2, describes the process for determining what threat to hunt for, evaluating it and then automating the further investigation.




Figure 2. Threat Hunting Loop

³ www.threathunting.net/sqrrl-archive

The threat hunting process is all about deciding what potential threat activity to look for, using tools to analyze the available data and teasing out patterns that could indicate a likely event. Each of these steps of the loop is unique to your organization, its infrastructure, the data available to the team and the tools at its disposal.

Create Hypothesis

Step one is to create the hypothesis. Did the attacker gain a foothold in the production web application? Could credentials be accidentally embedded in the packaged software? Is there an unknown, CPU-intensive process running on an important server? The sheer scope of potential hypotheses could grind any team progress to a halt.

Identifying and prioritizing the most at-risk infrastructure components requires an understanding of which systems are most vulnerable and their values to the business.⁴ By starting with a threat modeling process, an organization has an outline of priority systems that have a risk and are vulnerable to some set of attacks.

“At-risk infrastructure has one of four possible responses: attempt to mitigate the threat, eliminate the threat through infrastructure architecture, transfer the risk to a third party or just accept the risk.”

The threat hunting team needs to build a set of techniques to investigate and create a hypothesis of how those attacks would work and what artifacts are in the logs that need to be analyzed. Organizations with an offense-focused team, like a pen-test group or red team, have in-house experts who research and practice attacker techniques.

Others may need to rely on researching published materials on attack techniques to create new hypotheses. For example, the MITRE ATT&CK™ Framework is growing in popularity among researchers

⁴ Learn more about the threat modeling process in “How to Protect a Modern Web Application in AWS,” www.sans.org/reading-room/whitepapers/analyst/protect-modern-web-application-aws-38955, [Registration required.]

MITRE Enterprise ATTACK™ Framework

Persistence	Privilege Escalation	Defense Evasion	Credential Access	Discovery	Lateral Movement	Execution	Collection	Exfiltration	Command and Control
Image File Execution Options Injection			Forced Authentication	Network Share Discovery	Appletscript		Man in the Browser	Exfiltration Over Physical	Multi-Hop Proxy
Plist Modification			Hooking	System Time Discovery	Third-party Software		Browser Extensions	Medium	Domain Fronting
Valid Accounts			Password Filter DLL	Peripheral Device Discovery	Windows Remote Management		Video Capture	Exfiltration Over Command and Control Channel	Data Encoding
DLL Search Order Hijacking			LLMNR/NBT-NS Poisoning	Account Discovery	SSH Hijacking	LSASS Driver	Audio Capture	Remote File Copy	Remote File Copy
AppCert DLLs		Process Doppelganging	Securityd Memory	File and Directory Discovery	Distributed Component Object Model	Dynamic Data Exchange	Automated Collection	Scheduled Transfer	Multi-Stage Channels
Hooking		Muhta	Private Keys	System Information	Pass the Ticket	Muhta	Clipboard Data	Data Encrypted	Web Service
Startup Items		Hidden Files and Directories	Keychain	Discovery	Replication Through	Local Job Scheduling	Email Collection	Automated Exfiltration	Standard Non-Application Layer Protocol
Launch Daemon		Launchctl	Input Prompt	Security Software	Removable Media	Trap	Screen Capture	Exfiltration Over Other	Network Medium
Dylib Hijacking		Space after Filename	Bash History	Discovery	Windows Admin Shares	Source	Data Staged	Network Medium	Communication Through
Application Shimming		LC_MACH Hijacking	Two-Factor Authentication	System Network Connections	Remote Desktop Protocol	Launchctl	Input Capture	Exfiltration Over	Removable Media
AppKit DLLs		HISTCONTROL	Interception	Discovery	Pass the Hash	Space after Filename	Data from Network	Alternative Protocol	Multilayer Encryption
Web Shell		Hidden Users	Account Manipulation	System Owner/User	Exploitation of Vulnerability	Execution through Module	Shared Drive	Data Transfer Size Limits	Standard Application Layer Protocol
Service Registry Permissions Weakness		Clear Command History	Replication Through	Discovery	Shared Webroot	Load	Data from Local System	Data Compressed	Commonly Used Port
Scheduled Task		Gatekeeper Bypass	Removable Media	System Network Configuration	Logon Scripts	Regsvcs/Regasm	Data from Removable Media		Custom Cryptographic Protocol
New Service		Hidden Window	Input Capture	Discovery	Remote Services	Regsvr32			Standard Cryptographic Protocol
File System Permissions Weakness		Deobfuscate/Decode Files or Information	Network Sniffing	Application Window	Application Deployment	Execution through API			Custom Cryptographic Protocol
Path Interception		Trusted Developer Utilities	Credential Dumping	Discovery	Software	PowerShell			Data Obfuscation
Accessibility Features		Regsvcs/Regasm	Invoke Force	Network Service Scanning	Remote File Copy	Scripting			Custom Command and Control Protocol
Port Monitors		Exploitation of Vulnerability	Credentials in Files	Query Registry	Taint Shared Content				Connection Proxy
Screensaver				Remote System Discovery					Uncommonly Used Port
LSASS Driver		Extra Window Memory Injection		Permission Groups		Graphical User Interface			Multiband Communication
Browser Extensions		Access Token Manipulation		Discovery		Command-Line Interface			Fallback Channels
Local Job Scheduling		Bypass User Account Control		Process Discovery		Scheduled Task			
Re-opened Applications		Process Injection		System Service Discovery		Windows Management Instrumentation			
Rc common	SID-History Injection	Component Object Model				Trusted Developer Utilities			
Login Item	Subo	Hijacking				Service Execution			
LC_LOAD_DYLIB Addition	Setuid and Setgid	Installutil							
Launch Agent		Regsvr32							
Hidden Files and Directories		Code Signing							
.bash_profile and .bashrc		Modify Registry							
Trap		Component Firmware							
Launchctl		Redundant Access							

Figure 3. MITRE ATT&CK Framework⁵

and security companies (see Figure 3). Although not cloud-specific, the ATT&CK Framework provides a detailed explanation of the hows and whys of specific attacker techniques.

Specifically, the technique of gaining initial access by exploiting public-facing apps is relevant to the web app use case. ATT&CK describes the purpose of the technique, the types of platforms, potential mitigations and references to online reports. The information provided on this technique does not give us enough details to start hunting, but it does point to the Open Web Application Security Project (OWASP) Top 10, which is more relevant to the use case. More detail is noted in Figure 4.

When identifying the potential attacks against a web application, one of the best sources is the OWASP Top 10. The OWASP Top 10 is a documented explanation of the top security threats to web applications, detailing the attacker techniques, examples and potential ways to mitigate. The top threat in the OWASP Top 10 is an injection attack, or getting untrusted data sent to the interpreter and executed as part of a command or query. (See Figure 5.) In a SQL injection attack on a web server, the attacker provides unexpected values for the username or password to thwart the interpreter from retrieving the expected SQL values.

⁵ <https://attack.mitre.org/>

ENTERPRISE ▾

TECHNIQUES

All

Initial Access ▾

Drive-by Compromise

Exploit Public-Facing Application

External Remote Services

Hardware Additions

Replication Through Removable Media

Spearphishing Attachment

Spearphishing Link

Spearphishing via Service

Supply Chain Compromise

Trusted Relationship

Valid Accounts

Execution +

Persistence +

Privilege Escalation +

Defense Evasion +

Credential Access +

Discovery +

Lateral Movement +

Home > Techniques > Enterprise > Exploit Public-Facing Application

Exploit Public-Facing Application

The use of software, data, or commands to take advantage of a weakness in an Internet-facing computer system or program in order to cause unintended or unanticipated behavior. The weakness in the system can be a bug, a glitch, or a design vulnerability. These applications are often websites, but can include databases (like SQL) [1], standard services (like SMB [2] or SSH), and any other applications with Internet accessible open sockets, such as web servers and related services. [3] Depending on the flaw being exploited this may include Exploitation for Defense Evasion.

For websites and databases, the OWASP top 10 and CWE top 25 highlight the most common web-based vulnerabilities. [4] [5]

ID: T1190

Tactic: Initial Access

Platform: Linux, Windows, macOS

Data Sources: Packet capture, Web logs, Web application firewall logs, Application logs

Version: 1.1

Mitigations

Mitigation	Description
Application Isolation and Sandboxing	Application isolation will limit what other processes and system features the exploited target can access.
Exploit Protection	Web Application Firewalls may be used to limit exposure of applications to prevent exploit traffic from reaching the application.
Network Segmentation	Segment externally facing servers and services from the rest of the network with a DMZ or on separate hosting infrastructure.
Privileged Account Management	Use least privilege for service accounts will limit what permissions the exploited process gets on the rest of the system.
Update Software	Regularly scan externally facing systems for vulnerabilities and establish procedures to rapidly patch systems when critical vulnerabilities are discovered through scanning and through public disclosure.
Vulnerability Scanning	Regularly scan externally facing systems for vulnerabilities and establish procedures to rapidly patch systems when critical vulnerabilities are discovered through scanning and through public disclosure.

Examples

Figure 4. The Exploit Public-Facing Application Technique⁶

T10

OWASP Top 10

Application Security Risks – 2017

6

A1:2017-Injection

Injection flaws, such as SQL, NoSQL, OS, and LDAP injection, occur when untrusted data is sent to an interpreter as part of a command or query. The attacker's hostile data can trick the interpreter into executing unintended commands or accessing data without proper authorization.

Figure 5. Number One Threat in the OWASP Top 10⁷

⁶ Exploit Public-Facing Application," <https://attack.mitre.org/techniques/T1190/>

⁷ OWASP Top Ten Project, www.owasp.org/index.php/Category:OWASP_Top_Ten_Project

Other publications and researchers who track and describe attacker techniques include:

- Threat Post
- Threat Hunting Project
- AWS Security Bulletin
- (ISC)2 Cloud Security Report
- Summit Route
- Toni de la Fuente's running list of AWS Security Tools

The Cloud Security Alliance (CSA) publishes a report on top threats⁸ that focuses specifically on cloud services. The CSA also publishes an in-depth case study⁹ that walks through how those threats are carried out. Rhino Security is a pen-test company, but it publishes blogs and free tooling for cloud and containerization threats.

Investigate Via Tools and Techniques

Threat hunters go beyond the automated alerts from security products, past the strong indicators and into the squishy unknown. To do this, data must be collected, understood, analyzed and viewed comprehensively. Threat hunters must also pivot through different types of logs and explore unstructured or partially structured data.

The first hurdle can be the infrastructure itself. If the organization has dozens of unique operating system configurations, manually managed deployment or shared remote management, then logs and operational data will be highly variant, allowing real attacks to blend in. Let's look at another use case.

⁸Cloud Security Alliance, Top Threats to Cloud Computing: Egregious Eleven, <https://cloudsecurityalliance.org/artifacts/top-threats-to-cloud-computing-egregious-eleven/>

⁹Cloud Security Alliance, Top Threats to Cloud Computing: Deep Dive, <https://cloudsecurityalliance.org/artifacts/top-threats-to-cloud-computing-deep-dive/>




Figure 6. Overview of Amazon CloudWatch Log Collection

Use Case: Gathering SSH Connections

Leveraging infrastructure as code, it is possible to deploy production systems without administrators SSH'ing, except in cases of troubleshooting. Teams can easily pull logs from any system and into Amazon CloudWatch. See Figure 6.

To use the Amazon CloudWatch agent to pull SSH connection logs from Amazon EC2s and into the Amazon CloudWatch logging service, follow these steps:

1. Install the Amazon CloudWatch agent on an EC2.
2. Configure the Amazon CloudWatch agent to send SSH connections to a specific log group.
3. Set up Amazon CloudWatch alarms to monitor for invalid user attempts and repeated SSH disconnects.

The Ever-Changing Cloud Infrastructure

Cloud service elasticity can make it difficult to directly interrogate systems when the environment is continually growing and shrinking throughout a day. For example, let's say the web application is attacked at 10 p.m. with a SQL injection attack that triggers logs from the web application firewall (WAF). The next day at 9 a.m., the threat hunting team investigates to determine if the attack was successful.

Unfortunately, the VM has already been terminated by the cloud autoscaling engine. The threat hunting team needs to decide what data to collect from the elastic system, whether that data is readily available or needs to be pulled or pushed by additional systems, and how long to keep the data before aging it off. The threat hunter needs to account for the risk of those systems, the amount of data that might need to be stored and how quickly a team will evaluate the data. The following demonstrates an example.

Use Case: Post-Exploitation Detection

In a cloud environment of automation, once attackers gain access to the web application VM, they will want to use the MITRE ATT&CK tactic called Discover to find other services of interest, such as an accessible Amazon S3 bucket with the command **ListBuckets**. The web application we built has access to Amazon S3 buckets for configuration, but the IAM role does not allow listing of buckets. Automated systems likely already know the resources they need to interact with, so listing potential names is unnecessary. From the Amazon EC2 instance, listing buckets results in an error, as shown in Figure 7.

```
[ec2-user@ip-10-0-25-212 ~]$  
[ec2-user@ip-10-0-25-212 ~]$ aws s3api list-buckets  
  
An error occurred (AccessDenied) when calling the ListBuckets operation: Access Denied  
[ec2-user@ip-10-0-25-212 ~]$
```

Figure 7. A ListBuckets Error

AWS CloudTrail gathers and allows an analysis of Amazon Web Services (AWS) API requests. AWS CloudTrail, using the Amazon EC2 ID as the username, looks at the **ListBuckets** as an indicator. There is an **AccessDenied** error code, as shown in Figure 8.

Filter: User name		i-0b1515ec2d4b0b9df	Time range: Select time range	
Event time	User name	Event name	Resource type	Resource name
2019-09-14, 08:36:14 PM	i-0b1515ec2d4b0b9df	Decrypt		
2019-09-14, 08:35:37 PM	i-0b1515ec2d4b0b9df	ListBuckets		
AWS access key [REDACTED]		Event time 2019-09-14, 08:35:37 PM		
AWS region us-east-1		Read only true		
Error code AccessDenied		Request ID 244267745C55A876		
Event ID 396ef72d-8b25-4adc-a84a-7f0e4a09be3f		Source IP address 3.91.174.221		
Event name ListBuckets		User name i-0b1515ec2d4b0b9df		
Event source s3.amazonaws.com				

Figure 8. AccessDenied Error Code

Another option is to use the AWS Command Line Interface (CLI) to look for all commands from the Amazon EC2 in question:

```
aws cloudtrail lookup-events --lookup-attributes AttributeKey=Username,
AttributeValue=i-0b1515ec2d4b0b9df --query 'Events[].username:Username,
time:EventTime, event:EventName, eventId: EventId,resource:(Resources[0].
ResourceName)']}' --output table -- region us-east-1
```

Figure 9 shows sample results of AWS CloudTrail lookup-events.

LookupEvents				
event	eventid	resource	time	username
Decrypt	27bce37b-7db0-4567-8367-ee4f4f02ef39	None	1568507774.0	i-0b1515ec2d4b0b9df
ListBuckets	396ef72d-8b25-4adc-a84a-7f0e4a09be3f	None	1568507737.0	i-0b1515ec2d4b0b9df
ListBuckets	2579d4a9-e0b1-4cf0-b7a8-7f6edcab28ed	None	1568507736.0	i-0b1515ec2d4b0b9df
ListBuckets	aa9628dc-de9c-4818-8a40-dc22bc9dc846	None	1568507736.0	i-0b1515ec2d4b0b9df
ListBuckets	0b3c1151-7a61-4651-b91d-9f22a973cce5	None	1568507735.0	i-0b1515ec2d4b0b9df
ListBuckets	5a1384c4-b77d-46e0-8c6d-4486a15ddb37	None	1568507365.0	i-0b1515ec2d4b0b9df
ListBuckets	8f8158ff-b837-4bba-a413-43ebcc65107b	None	1568507363.0	i-0b1515ec2d4b0b9df
ListBuckets	13485b09-a4e8-4e62-aec1-c4d5982e86b3	None	1568507362.0	i-0b1515ec2d4b0b9df
ListBuckets	1ef3785c-c2cb-4ee0-bb60-807c8e00b9b8	None	1568507361.0	i-0b1515ec2d4b0b9df
ListBuckets	373507ce-1331-4682-81b7-313a260bcd7e	None	1568507309.0	i-0b1515ec2d4b0b9df

Figure 9. Table Output of AWS CloudTrail lookup-events Command

Each event has a unique event ID. Figure 10 shows the details for a specific event ID from the table shown in Figure 9. Here, we use a Linux application, JQ, to carve up JSON on the command line. This command shows the details of this particular AWS CloudTrail Event. JQ is an excellent tool for filtering, carving and formatting the JSON data in logs.

```
cybergoof> aws cloudtrail lookup-events --lookup-attribute AttributeKey=EventId,AttributeValue=396ef72d-8b25-4adc-a84a-7f0e4a09be3f --query "Events[0].CloudTrailEvent" --region us-east-1 --output text | jq -r '.'
{
  "eventVersion": "1.05",
  "userIdentity": {
    "type": "AssumedRole",
```

Figure 10. JSON Output of AWS CloudTrail lookup-events

Uncover New Patterns and Apply Learned Lessons

Gathering data, running analytics and identifying the anomalies give the threat hunter unique insights into evaluating attack techniques and analyzing infrastructure systems. The team should become part of the threat modeling processes, helping the architecture and operations teams identify the cloud infrastructure that needs to be secured and evaluated. Changes such as improved monitoring, reduced chaotic deployments and better segmentation of infrastructure can all make threat hunting easier without losing operational capabilities.

Once threat hunters understand the challenges, they can start gathering detailed knowledge of potential threats, and the architecture and infrastructure management teams can support the threat hunters. It is time to begin collecting and analyzing the data needed to discover the attackers.

Inform with Data and Analytics

It is critical to get the right data into the right place for analysis. The data itself might need to be evaluated, enriched and prepared for analysis using scripts, tools or built-in cloud services.

Gathering the Data

The threat hunting team has to strike the right balance of how much data to capture. Requiring all the data from all the things increases costs, adds to the overhead of managing the data and increases the time and effort to sift through and analyze the enormous amounts of data. On the other hand, not having enough data will keep the threat hunters in the dark. First, identify any logs that are already being collected or are easy to obtain organically. AWS makes it easy to collect VPC logs showing data connections in and out of the VPC, API calls with AWS CloudTrail and Amazon S3 access logs, among others.

Then, using the attacker techniques, the team will focus on identifying the gaps in information and how to retrieve it. Most missing data is likely from applications or the host environment itself. Let's revisit the web application use case.

Web Application Use Case

For the web application use case, the VM itself has a wealth of information that could be of interest. Mainstream web servers generate standard logs that are stored on the VM. They also can be customized to generate more or fewer logs, or with changes to the format or location, and potentially compressed

for transfer. Connection logs, for example, contain every HTTP request to the web server. Regularly managed web applications have a lot of the same connections. However, in a path traversal attack,¹⁰ the path could contain unique path calls that are attempts to get access to files on the web server.

After installing the Amazon CloudWatch agent, configure the Amazon CloudWatch configuration file to pull the Nginx access log **/var/log/nginx/access.log**. See Figure 11.

```
[/var/log/nginx]
datetime_format = %b %d %H:%M:%S
file = /var/log/nginx/access.log
buffer_duration = 5000
log_stream_name = {instance_id}
initial_position = start_of_file
log_group_name = /var/log/nginx
```

Figure 11. Amazon CloudWatch Logs Configuration File

The Nginx connection logs are now stored in the **/var/log/nginx** log group, accessible from Amazon CloudWatch Logs. See Figure 12.

Create Metric Filter	Actions ▾
Filter:	Log Group Name Prefix x
Log Groups	Insights
☐ /var/log/messages	Explore
☐ /var/log/nginx	Explore

Figure 12. Nginx Connection Logs

passwd	
Time (UTC +00:00)	Message
2019-09-15	
No older events found for the selected filter and date range. Adjust the date range or clear filter .	
▶ 23:56:18	173.69.145.155 - - [15/Sep/2019:22:41:41 +0000] "GET /?file=../../etc/passwd HTTP/1.1" 200 1378 "-" "Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/74.0.3683.103 Safari/537.36"
▶ 23:56:18	173.69.145.155 - - [15/Sep/2019:23:48:21 +0000] "GET /?file=../../etc/passwd HTTP/1.1" 304 0 "-" "Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/74.0.3683.103 Safari/537.36"
▶ 23:56:18	173.69.145.155 - - [15/Sep/2019:23:48:28 +0000] "GET /?file=../../etc/passwd HTTP/1.1" 304 0 "-" "Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/74.0.3683.103 Safari/537.36"
▶ 23:56:18	173.69.145.155 - - [15/Sep/2019:23:49:11 +0000] "GET /passwd HTTP/1.1" 404 3696 "-" "Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/74.0.3683.103 Safari/537.36"

Figure 13. Quick Search for passwd

¹⁰ www.owasp.org/index.php/Path_Traversal

Opening up the log group, it's possible to search for a string, as shown in Figure 13.

This is an easy search. AWS provides an advanced query service called Amazon CloudWatch Logs Insights. Using a custom query language, we can search across all hosts for a regex of passwd, etc or ../ as shown in Figure 14. Note that / is a special character in regular expression (regex), so it has to be escaped with \.

Figure 15 shows the results of the query. Once the data is gathered, the data retention life cycle rule is applied and data is accessible, it's time to figure out how to make the data more useful to the threat hunters by enriching the data.

```
fields @timestamp, @message
| sort @timestamp desc
| limit 20
| filter @message like /passwd|etc|..\|/|
```

Figure 14. Query Amazon CloudWatch Logs Insights




Figure 15. Query Results

Enriching the Data

When threat hunting, the data needs to tell a complex and complete story with multiple characters, settings and subplots. If a single log could tell the story, then a security product would quickly alert the SOC. Threat hunters are looking for more subtle anomalies in the data that look unique mainly because of the way an infrastructure is architected and operated. An attachment in the email is easily scanned and compared to a known list of malware. However, it's harder to identify a nefarious remote desktop connection compared to a legitimate one. One easy way to bring data to life is to automatically evaluate the data and tag it, add metadata or enhance the data itself.

Separate Security Account:

It is good to gather and protect any logs from accidental or purposeful deletion. One recommendation is to use AWS Organizations to create a separate security organization (org) and to automatically move logs from the production org to the security org, where it can be protected and available to only the security or designated teams.

Web Application Use Case

There are several ways to automate the analysis and tagging or enriching the data. For logs collected by Amazon CloudWatch, such as Nginx connection logs, leveraging the alarms, metrics and dashboards works well. An Amazon CloudWatch Metric Filter will search for some specific patterns and create a metric count when that pattern shows up in the logs. An Amazon CloudWatch metric can generate an alarm, which can send an email or notify an AWS Lambda function. The AWS Lambda function can take action, such as copying the concerning data over to an Amazon S3 bucket for further analysis.

In the Amazon EC2 Role use case, the victim EC2 can perform S3 bucket reads. Let's say there are 50 EC2 instances in the account; that would be too much data to analyze. However, if the EC2 reads a different S3 bucket than it has ever read before, that is a new activity. You should tag those reads.

Analyzing the Data

Once the data has been gathered, enriched and tagged, the threat hunting team starts evaluating the data to identify anomalous behaviors against the hypothetical attack techniques. The threat hunting

team must be able to evaluate anomalies and quickly determine if they warrant an investigation or not, so the data must be easy to search, correlate and report. Various scripting tools and analytic platforms can provide threat hunters with raw log data to sift through. Comprehensive analytic platforms can also be utilized to help speed up analysis, and provide reporting services for sharing and collaboration among teams.

The next sections dive into options for analytic tools to bring into the environment to take threat hunting to the next level.

Tools for Analysis

Threat hunters can bring a wide range of tools to bear to analyze complex datasets from multiple sources, from scripts parsing raw data, to a full SIEM system that provides ad hoc and complex searching, reporting and investigations. The decision is usually about setup complexity, cost and the need to scale as the team grows. AWS provides several services that can be used and chained together to scripts and analytics.

Analyzing Logs Directly

Amazon CloudWatch is the core service for monitoring an AWS environment, because it is easy to get up and running and providing basic metrics, alarming and dashboards. As was previously discussed, Amazon CloudWatch and AWS CloudTrail can be used together to interact directly with collected data. AWS offers methods of exporting Amazon CloudWatch logs, collected from custom applications to Amazon S3, AWS Lambda or Amazon Elasticsearch Service (see Figure 16).




Figure 16. Exporting Amazon CloudWatch Logs

AWS provides another service called Amazon Athena, which runs SQL queries against data in an Amazon S3 bucket (see Figure 17). Customers build virtual tables that organize and format the underlining log data inside the bucket objects. It takes time to ensure that data is formatted and managed.

Amazon GuardDuty is a managed service that is evaluating a growing number of findings that detect adversary behaviors and alerting the customer. Amazon GuardDuty evaluates potential behaviors by analyzing Amazon VPC Flow Logs. A similar real-time VPC flow logs analysis engine can be created using AWS Lambda, Amazon Kinesis, Amazon S3, Amazon Athena and Amazon QuickSight.

SIEMs in the Cloud

As a threat hunting team starts to build a corpus of analytics that it wants to run repeatedly, or as its investigating, monitoring and reporting needs become more comprehensive, a full SIEM is likely of interest. Several cloud-specific services, as well as traditional on-premises SIEMs, work with cloud infrastructure.




Figure 17. Amazon Athena Dashboard

The threat hunting team should focus on developing and managing a tactical SIEM, which could be different from the SIEM a SOC might use. The tactical SIEM will likely have unstructured data, a shorter retention policy than the SOC's SIEM, and the ability to easily determine what the infrastructure looked like in the recent past. In the cloud, good data management strategy should be implemented to be cost-effective, with pay-per-usage pricing. Generally, free or open source solutions tend to take more time and expertise to set up and maintain, but they are more customizable and cost little or nothing. Commercial solutions may cost more, but may come with better support, easy access to purpose-built connectors and more reporting options.

Elasticsearch, a favorite of the open source community, boasts a significant user base and supports plug-ins for data importing, translating and easy displaying with the Kibana application. AWS provides a managed Amazon Elasticsearch Service to make it easy to set up and run the search engine without having to do all the management heavy lifting. The company behind Elasticsearch, Elastic, has released a new app called the Elastic SIEM that is more focused on the security operations. Other products, such as ones from Sumo Logic and Splunk, also integrate directly with AWS and provide even richer and more full-featured analytic platforms.

After the tactical SIEM is stood up; the data is gathered, translated and enriched; and mechanisms for analytics and reporting are in place, the threat hunting team will start to discover repeated steps, analytics or actions. An emerging service that integrates with the SIEM, called Security Orchestration, Automation and Response (SOAR), can be helpful there.

Soaring with SOAR

Threat hunting is all about proactive analysis of data to detect the anomalous behavior that is undetectable by the security products. As the threat hunting team's analytics become more sophisticated, it may begin developing a set of repeatable analytics, enrichments or data gathering steps. If it's repeatable and articulate, it can be automated. A SOAR leverages the data storage and enrichment of the SIEM, understands basic rules of infrastructure integration and allows the easy buildout of playbooks to automate a course of action.

In the web application use case, if there are several failed SQL injection attempts, the final attempt could signify the last failure before success. The process of information from that host at that time would be of interest. A SOAR could be used to identify that ultimate SQL injection failure, tag it and then also tag the process log information from that time. The next step in the playbook could be to move those logs into a separate Amazon S3 bucket for more accessible analysis. The process logs by themselves could then be enriched by validating with a malware signature API to identify whether the process is known good or

not. Gathering potential logs to analyze and automating the enriching processes when necessary could save threat hunters tedious and repetitive work. It could also help provide quicker triage. The SIEM with a SOAR could significantly improve speed to analysis.

Taking the playbook a step further, it's possible to use data pushed to the SIEM and SOAR, such as the SQL injection detection logs from the WAF, and initiate an action. Rather than always pull the process list on an hourly basis, the SIEM could execute host-based tools, such as OSQuery, to reach out to the suspect web server and pull the process list in near real time. This automated response action allows the team to limit what passive data has to be managed, and makes it easier to correlate the process logs returned with the suspicious SQL injection attacks.

In the Amazon EC2 use case, the SIEM/SOAR could review the READs from an EC2 to an Amazon S3 bucket and detect a first-time READ to an S3 bucket. The SOAR playbook executes a host agent such as OSQuery or uses AWS services such as Amazon Inspector and AWS Systems Manager to interact directly with that EC2 to pull fresh process information and kick off a scan with Amazon Inspector. It then gathers all these reports and provides them in a single artifact bucket for the security analysts, creating a high-priority message in the corporate chat system or sending out SMS alerts to on-call personnel.

Some of the more sophisticated SOARs, such as Splunk's Phantom, also allow for the detection of cascading anomaly triggers that can perform automated remediations—taking our use cases together to build a sophisticated SOAR playbook.

“As the threat hunting team’s analytics become more sophisticated, it may begin developing a set of repeatable analytics, enrichments or data gathering steps. If it’s repeatable and articulate, it can be automated.”

SOAR Playbook Use Case

The attacker performs several SQL injection attacks against a particular EC2. The SOAR kicks off a process listing and tags all logs from that EC2 with a unique identifier. One of those logs with the unique identifier specifies a failed Amazon S3 bucket listing attempt. In an automated system, the bucket is

known, and a listing is unlikely to be normal. The SOAR identifies that this failed bucket listing happened on an EC2 that is being triaged. Because the organization is using auto-scaling, the SOAR notifies the auto-scaling system to deregister the EC2 (i.e., pull that EC2 out of service but keep it running). The SOAR playbook waits for the deregistering to finish, then removes all security groups except triage, and the triage group effectively isolates the EC2 from all other systems.

Conclusion

We are in the early days of threat hunting, specifically in cloud environments. Organizations are moving away from traditional server-based infrastructure into serverless, event-driven architectures that rely on native cloud services. Threat hunters will adapt their processes, tools and techniques to identify and neutralize the threats in this new infrastructure landscape.

Threat hunting is critical to finding the advanced attacker techniques that have escaped the detection of deployed security products. The threat hunting process requires constant learning about attacker techniques and your organization's attack surface. Proper strategy ensures the right data is collected, enriched and available to the tools the threat hunting team uses to tease out suspicious anomalies from the vast and ever-changing infrastructure. Your threat hunting process is always growing and adapting to new learnings, increasing experience and the changing threat landscape.

About the Author

During his 25+ years of experience, Shaun has spent equal parts in security engineer and operations as well as software development. With extensive experience within the Department of Defense, Shaun was the Technical Director of the Red and Blue operations teams, a researcher of advanced host analytics, and ran a threat intelligence focused open source platform based on MITRE ATT&CK. Previously, he was a consultant with H&A Security Solutions, focusing on analytic development, DevOps support, and security automation tooling. Shaun has authored the brand new SEC541: Cloud Monitoring and Threat Hunting and can be found teaching SEC545: Cloud Security Architecture and Operations on a regular basis.

Chapter 20: Solution Guidance for Application Security in AWS



Nathan Getty

SANS Analyst

"In this chapter, I cover a wide variety of topics in regard to application security in the AWS Cloud. This includes things like policies and standards to implement that aid the organizations application deployment, coding standards and the software development life cycle (SDLC), specific ways to detect flaws in your codebase (inline scanning/out of band scanning), open source tools that can statically and dynamically check the posture of our codebase/application stack (linters), and key considerations to evaluate in the AWS cloud environment.

This chapter targets security analysts who would like to broaden their depth in the AWS application security plane. There may be some topics or acronyms that are new, but a few hours of independent research should bring you up to speed. As the title of this series is "JumpStart," this chapter should be great starting point for security analysts and application security engineers who plan to broaden their knowledge in the AWS AppSec world."

Introduction

As organizations begin to transition their applications into cloud environments, security teams must provide application security support and insight during the process.

Today's applications are updated more frequently, and regular release cycles are giving way to more rapid incremental releases. Application development continues to evolve to support a more dynamic release schedule. In response, information security teams must be included in the development process if they are to provide support to development teams. Because organizations plan to deploy applications as soon as they are approved for production, your organization's security team should not be the roadblock.

Because development teams release applications faster than they can be reviewed, it is critical to integrate the skills and guidance of the security team into the development model. Whether the application code is deployed on premises or in a cloud environment, automated security tools provide the information security team with visibility into code as it moves through the developer pipeline. This visibility provides more assurance that security will not be compromised.

This process allows the development teams to remain informed of security concerns for their application as it moves through the pipeline. By embedding security within the build process, your organization can build a strong relationship between the security and development teams. By fostering and developing this relationship, developers and security professionals can work in tandem to deliver secure, timely applications.

According to Forbes, nearly three-quarters of companies are planning to move to a fully software-defined data center within two years. Almost half of businesses are delaying cloud deployment due to a cybersecurity skills gap.¹ This paper seeks to give you a better idea of what your organization needs to successfully plan and execute a secure application transition to, or deployment in, an AWS environment.² We discuss how security teams can best support application development teams, what options you have as a security professional for this support, and how best to guide your development teams as they transition workflows to AWS.

¹"2017 State of Cloud Adoption and Security," www.forbes.com/sites/louiscolumbus/2017/04/23/2017-state-of-cloud-adoption-and-security

²This paper mentions product names to provide real-life examples of how firewall tools can be used. The use of these examples is not an endorsement of any product.

Understanding Your Needs

Historically, application development and security teams did not always work closely together. But given the adoption of rapid release cycles and the transition to cloud services, these teams must build a working relationship that effectively supports rapid deployment of secure applications. How can they do that while best using existing tools and processes in the cloud environment?

1. Understand the applications deployed in your organization.

Security analysts need to be knowledgeable about the applications being deployed, at least to the extent of being aware of their primary purpose and target audience. When they understand the application, the underlying code, and for whom the application is designed, they can run threat modeling assessments and plan accordingly. They can make remediation decisions with confidence, bring attention to specific security vulnerabilities, identify which vulnerabilities and risks are acceptable, and provide feedback to the development team. Encouraging security teams to work closely with development teams and speak their language will build a strong, mutually beneficial relationship.

2. Understand application deployment methods within AWS.

Applications can be deployed through any one of several channels or tools. Knowledge of the tools available to development teams can help information security teams define best security practices within those tools and ease incident response or critical changes to the applications. Through awareness of the underlying development process, an organization can be assured that quality information regarding security concerns is being communicated to the development teams.

3. Understand what options and responsibilities you have in AWS as you prepare for securing the application delivery.

The AWS cloud environment gives organizations access to a large developmental toolset in the form of services that include a number of capabilities. Not every service will be a good fit for your organization, so development and security teams should plan ahead and identify which services they will need to use for their application delivery and the security.

AWS offers various platforms for setting up such services. For example, AWS offers serverless services, which means your organization is not responsible for operating or maintaining the underlying infrastructure. Although AWS takes full responsibility for operating the hardware, networking and patch management of the underlying infrastructure, responsibility for the security of any application built on the platform lies completely with the organization.

Implementation Options in AWS

AWS offers a number of services and options as well as access to third-party services for secure application development and rapid release cycles.

Cloud-Native Services

When applications and security tools work harmoniously, future problems (and the need to fix them) can be avoided more easily. Fortunately, AWS-native services are built to work well with each other. Leveraging native services can ease the speed of deployment and integration of application security tools. AWS Marketplace contains a collection of ready-to-deploy infrastructure components your organization can deploy directly into their Amazon VPC (Virtual Private Cloud). AWS Marketplace offers a variety of software including, but not limited to, operating systems, network and business intelligence tools, machine learning software, security software and development suites.

The ability to find, test, deploy and validate software through AWS Marketplace helps organizations identify which applications work for them, which allows them to procure and deploy solutions much faster than when having to spend time engaging with a variety of vendors. (Although deploying AWS Marketplace products can be quick and fast, you should still engage with your organization's software onboarding team before deploying new solutions within your environment; your organization may have certain software onboarding procedures even when it comes to native AWS services.) Leveraging native services also has the added benefit of pricing consolidation. Because AWS services are billed to your account with detailed information, organizations can use native services to view all of their AWS costs within a single, detailed page.

Open Source and Custom Solutions

Native services offer direct benefit to your organization, but there may be situations where you prefer custom or open source software (OSS) applications. OSS and custom tools can be leveraged within AWS as long as they are compatible with AWS infrastructure (Microsoft Windows- or Linux-based platforms). For example, it is possible to run custom or OSS solutions on Amazon EC2 (Elastic Cloud Compute). The key difference with EC2 (versus native service) is that your organization inherits the full responsibility for any underlying infrastructure. Your organization is responsible for patch management and any security solutions required for the infrastructure (firewall, intrusion detection and other security tools). Refer to the AWS Shared Responsibility Model³ for more information.

³ AWS Shared Responsibility Model, <https://aws.amazon.com/compliance/shared-responsibility-model/>

Consulting Partner Private Offers

Customers can also engage through Consulting Partner Private Offers (CPPO) to work directly with trusted advisors to select and configure Application Security solutions from AWS Marketplace. As organizations build out their cloud and cloud security strategy and plan, they may want to consider working with partners to accelerate their efforts or fill any gaps in knowledge or resources that are identified. All consulting partners may extend AWS Marketplace third-party solutions directly to customers through CPPO.⁴ Not every organization will be able to find resources with deep cloud experience. Even experienced cloud technologists may have experience only with specific industries or cloud vendors. A requirements document could be helpful when approaching prospective consultants.

Needs and Capabilities: The Business Case for Application Security in the Cloud

The benefits of putting applications in the cloud must be balanced by the organization's ability to secure them.

Application Security

The need: Conducting application security assessments and reducing vulnerabilities within the AWS environment

Capabilities

- Increased visibility within the development process and application stack
- Reduced risk and vulnerabilities in the applications before they are deployed
- Automated security assessments with actionable remediation
- A relationship with the development teams

⁴ AWS Marketplace Channel Programs, <https://aws.amazon.com/marketplace/partners/channel-programs>

General AWS Web Application Security Considerations

Regardless of the technology or cloud vendor selected, some general business, technical and operational considerations are associated with implementing application security in the cloud. The following sections highlight many of these considerations.

Business Considerations

	Consideration	Details
	Policies and standards	Organizations must understand their current software development life cycle (SDLC) policy and how it may be affected by a move to a cloud environment. An SDLC policy describes the various stages of application deployment and delivery. These underlying methodologies do not change when moving to a cloud environment but the processes and procedures for application code review, application building, delivery and analysis probably will. Anticipating what changes to the SDLC will be triggered by transitioning to AWS will allow organizations to adopt an SDLC that not only fits the cloud model, but also has tangible benefits for an organization's application delivery within the cloud. Planning and making these changes first will save your organization time should a policy need to be redefined in the future. Organizations should determine the acceptable level of risk for their application(s). Although it would be nice if we could deliver applications without errors or exploitable weaknesses, such a scenario is unfortunately unrealistic. Developers have to release applications within the timelines demanded by their sprints, and they often lack sufficient resources to explore and address all security aspects of their application in the available time. If an organization deploys an application with little or no security validation, it is exposed to a greater risk that the application could be exploited. Organizations must plan ahead and define an acceptable threshold for vulnerabilities within a production-class application. For example: Organization X ships releases for its Acme web app every two weeks. It runs security tests each time the application is built. Its policy states that if those tests find that the application build contains more than three high-risk vulnerabilities or greater than zero critical risk vulnerabilities, Organization X will block application delivery until the issues have been addressed and corrected.
	Licensing options	While AWS operates under the “pay what you use” model, many third-party vendors allow customers to deploy products directly on AWS's infrastructure. Leveraging third-party applications and tools can quickly increase licensing costs for your organization. Take precautions when deploying third-party applications and tools on AWS infrastructure, because your organization will incur both AWS infrastructure usage and software licensing costs. Licensing costs can be charged in a few different ways. They may be billed to the organization on an annual basis or perhaps by the hour. Understanding and planning for expected licensing costs will ensure you are not caught off guard by large invoices from AWS.

Technical Considerations

	Consideration	Details
	Technology deployment	Organizations should plan ways to implement their application security in a repeatable, consumable manner. Security teams can provide guidance in this matter in a variety of ways. Within AWS, applications can be deployed through a fully automated “pipeline”; alternatively, they can be deployed in an ad hoc fashion. An organization would be wise to create small, repeatable security tests as part of the deployment process, and to continuously refine those tests as the application matures. Understanding how your organization deploys its applications will allow the security teams to create and deploy effective security tests that align with the developers’ deployment plan. Organizations need to decide if they will allow OSS or unsupported technologies. While it’s true that an open source application allows insightful visibility into the application’s security, it’s also true that open source projects do not come with the luxury of customer support or SLA. If you plan to use open source technology for critical tasks or security assurance, you will need to ensure you have a proper plan in case the tool stops working at some point. On the other hand, OSS tools offer some unique opportunities. Organizations can take advantage of free open source tools and, as their needs outgrow the capabilities, modularity or support level provided by the OSS tooling, they can transition to more professional offerings.

	Consideration	Details
	Application stack	AWS Marketplace offers many tools for securing your organization’s applications. Leverage any available open source testing software to get used to integrating security tools into your application development process (and save costs). Static analysis tools (linters) allow you to check your code for programming errors, bugs, stylistic problems and suspicious constructs. Each programming language has its own set of linters, most of which can be installed directly within your developers’ preferred integrated development environment (IDE). Having developers use a linter within their IDE saves time in the development process by catching the errors before the application code is pushed. Catching these issues before the application is deployed makes it easier to mitigate them after deployment. Organizations should also consider their application stack and what corresponding Static Analysis Security Testing (SAST) tools might best fit their deployment pipeline. While linters check for bugs, syntactical errors, programmatic errors and code nuances, the purpose of SAST tools is to identify security issues in the application source code (versus during compilation or runtime). As with linters, each language has its own set of SAST tools, so your organization needs to understand the application code being implemented and what the information security teams will need to deploy to validate the codebase.

Technical Considerations Continued

	Consideration	Details
	Pre-deployment security (inline)	The largest challenge of inline scanning is the time it takes scans to complete. If your organization needs to deploy an application change, your security test should not require a long time to run. Imagine making a small configuration change to your organization's application. You push your code to the development pipeline, and now you have to wait 30 minutes for the security tools to scan your changes. Developers can push these changes many times a day, so waiting for these scans can be frustrating. We recommend that inline scans should not take longer than five minutes (depending on the size of the codebase). Your organization might also want to consider scanning only the changes to the code from the last push (delta scan). This method saves time but may be better suited to more mature organizations. It also makes sense to occasionally scan the entire codebase outside of the pipeline (out-of-band scans). We advise that organizations take small, repeatable, incremental steps in deploying inline scans for application pipelines. It's a good idea for your security team to have its own source code repository where it stores its tests. After a test has been created and validated, it can be stored in the repository. Once the code is in this repository, it may be shared with the developers, and they can include them within their development pipeline. You can work with the developers to ensure that the latest copy of the security test is always referred to when inline scanning. This procedure allows the security team to update the test as it sees fit. Because the development team has the latest copy of the test always being pulled into the pipeline, there should be no additional work when the security tests are updated. Leveraging this approach allows you to continuously test applications, update the tests and keep track of what exactly was changed via revision control.
	Post-deployment security (out of band)	Organizations will need to decide when to implement post-deployment security scanning. We mentioned out-of-band scanning earlier: If scans take too long to complete, they can be scheduled after the application has been deployed. Full scans by Dynamic Application Security Testing (DAST) tools can take hours to run, depending on the application size and scope of the scan. The following are examples of tools that should be run outside of the deployment pipeline: • Infrastructure scans—These can take a long time depending on the scope of the resources and security checks the scan performs. • Dynamic application security scans—These require the environment to already be up and running. Like infrastructure scans, these scans can take some time to complete, depending on the organization's scanning scope. • Full web application security scans—Depending on the parameters of the test (credentialed/no-credential/spider/full active scan) and the size of the application, this scan can take a long time to run and should not be used inline. Organizations will need to decide what is necessary to test and ensure application security for applications that have already been deployed. Solutions such as infrastructure security scanning, WAF implementation and DDoS protection should be evaluated.

Operational Considerations

	Consideration	Details
	Processes and procedures	Organizations may need to create or modify processes and procedures for security web applications in AWS. While some existing processes and procedures may work without modification, hosting applications in AWS means different methods of application delivery. Organizations may want to start to include developers and key individuals involved with application delivery in meetings and discussions about application security testing. Security teams might also want to sit in on development meetings and inform discussions when application security concerns arise.
	Resources and deployment synergy	Security in AWS and the applications deployed within the cloud will take dedicated resources to ensure that the proper policies and procedures are followed. Organizations must be cognizant that resources will need to be dedicated in such an effort, and they should plan accordingly. Organizations should consider which approach they would like to take with their cloud application security and the level of responsibility for each team involved within the process. Development and security teams within your organization need to take responsibility for the security and integrity of the application.

AWS Implementation Considerations

Application Security

	Consideration	Details
	Cloud context support	Application deployment leverages many ephemeral resources that support application delivery. Catalog all possible resources used within the deployment process for identifying any issues. Evaluate: • The additional cloud context (tags or image IDs, other possible ephemeral resources) captured within the development processes (phoenix servers, artifacts and the like) • Logging and cataloging of the cloud resources for traceability and troubleshooting
	Deployment	Deployment methods for security tools within AWS can vary depending on the development pipeline. Organizations should deploy these tools within the context of the development pipeline. Evaluate: • Installation and initial configuration for tools • Possible use of professional services to aid or accelerate tool deployment • Programming tools and languages used in the applications and their corresponding DAST/SAST tools • The availability of managed or SaaS components or preconfigured appliances from AWS Marketplace • Leveraging AWS-native services for security implementation
	Integration	Integration of application security tools into current processes/procedures ensures security teams can respond to risks. Integrating application security tools into the development pipeline allows for visibility, deployment and management. It also provides ease of use for security and development teams. Evaluate: • The development pipeline process and how to embed security tools and scans inline within a reasonable time • Tools that integrate with current security solutions (SIEM, SOAR, IT service management) • API support (REST APIs available, SOAP APIs available, other available programmatic APIs) • Use of custom plugins or integrations • Integrations with native AWS services

Making the Choice

To summarize, the key considerations for implementing application security in AWS are:

- Cloud context
- Deployment
- Integration
- Configuration and iteration
- Reporting

Evaluate Your Organization's Current Deployment Process

There are many ways to deploy applications with AWS, and many methods with which to build out your deployment pipeline. When defining your proof of concept, include significant members of the application deployment team and ensure you understand their method of deployment infrastructure (Amazon EC2, Amazon ECS, serverless) and deployment pipelines (AWS CodePipeline, Jenkins, other deployment tools). Once your organization has a strong understanding of the deployment process, it can begin to evaluate its needs and considerations for security tools. Define a proof of concept that meets both your organization's considerations and the developers' current deployment process.

Define a Plan and Implement

By defining and understanding its cloud architecture, risk profile, business requirements and available resources as well as all the possible deployment methods within AWS, an organization should have a clear idea of its road map for application security protection. Understand that defining application security that meets all the discussed considerations is nearly impossible, so define and use what works best for your organization.

The best course of action is to define a proof-of-concept plan based on the considerations and implementation options. Ensure that your organization's development team is included in this process, because they will have a very strong understanding of the application and which security concerns to note. Once you have planned, developed and validated your POC, development and security teams can start defining a repeatable process for integrating app security within the development process. In this

stage, your organization should work with the development team to identify the team's current security issues and how the developed POC will help secure the application and reduce the application's risk to meet the organization's risk threshold.




Figure 1. The DevOps Life Cycle

Conclusion

Application security is a crucial step for organizations' cloud security strategy. Having a defined plan and integrating security within the development process allows for greater visibility within the application delivery process, visibility into the security stance of the application and a defined remediation process for application security vulnerabilities.

Work with the development team through each stage of the DevOps life cycle (see Figure 1). Plan with the developers, join meetings when they develop and discuss their applications, ask the developer team for help when writing security tests in the verification stage, add out-of-band security (WAF protections, EDR solutions, DDoS protections and the like) in the release stage, and constantly monitor the security state of the application through your infrastructure monitoring and log analysis. Security tools and checks can be applied to many stages of the development process.

Keep in mind that this process should always be repeatable and easy to use. Start small and build from there. To get started today, consider an evaluation of some of the solutions readily available via the AWS Marketplace. You may also consider leveraging a SaaS solution to jump-start your organization's journey into AWS application security.

About the Author

Nathan Getty holds the GWAPT and GCIA certifications, and he recently won the SANS Cyber Defense NetWars competition, a defense-focused challenge that tests the ability to solve problems and secure systems from compromise. Nathan currently works for one of the world's leading food delivery companies. In his organization, he focuses on cloud security, including AWS onboarding, and developing best security practices and general security/ cloud insights. Nathan also focuses on driving DevOps methodologies in the company's security program, implementing continuous delivery platforms to allow smoother development and improvement of internal security applications.

Chapter 21: Solution Guidance for Cloud-Based Firewalls in AWS



Brian Russell

SANS Analyst

"This chapter provides an overview of the implementation options for firewalls and threat prevention in AWS. In this chapter, I discuss the needs and capabilities associated with these products and review the key technical and operational considerations that must be considered when planning an implementation. I review options such as bring your own license (BYOL), managed firewall and virtual firewall appliances available in the AWS Marketplace.

The target audience includes cyber security engineers and cloud engineers responsible for integrating a secure AWS environment. This chapter concludes with a simple method for performing an analysis of alternatives that hopefully aids in the decision-making process for your unique environment."

Introduction

Firewalls have evolved from providing simple packet filtering based on port and protocol combinations. Today's cloud-based firewalls are virtualized in the cloud and provide rich features such as application-based filtering, microsegmentation, encrypted traffic inspection and DNS security. Cloud-based firewalls are becoming true security platforms that incorporate intrusion prevention and detection features and threat prevention services that allow organizations to stay protected against both known and unknown malware.

This guide examines options for implementing firewalls within Amazon Web Services (AWS). It examines the needs and capabilities associated with today's firewall and threat prevention services and details general, technical and operational considerations when choosing these products. The guide concludes by examining AWS-specific considerations and recommending a plan of action for organizations considering the purchase of cloud-based firewalls. Before we begin, Table 1 provides definitions of key firewall-related terms.

The considerations in this guide are designed to inform a systematic evaluation strategy for choosing the optimal firewall for your requirements. An evaluation strategy should be based on an organization's specific needs and implementation requirements. The evaluation should consider the capabilities of the native AWS firewall offerings and then incorporate a review and comparison of AWS Marketplace offerings. Finally, following the simple "Analysis of Alternatives" detailed in the "Making the Choice" section of this paper will assist you in making the right decision for your organization.¹

¹ "Analysis of Alternatives," https://en.wikipedia.org/wiki/Analysis_of_Alternatives

Table 1. Key Terminology in This Guide	
Terms	Descriptions
Network Firewall	Network security device used to monitor incoming traffic and block unauthorized traffic. Commonly, a set of rules is defined for ingress and egress traffic. Only authorized traffic is allowed into and out of the network. Rules are typically set up based on IP address and port combinations.
Web Application Firewall	An HTTP application-specific firewall used to protect an application's back-end servers from attacks such as cross-site scripting and SQL injection. A set of rules governing the format and content of HTTP messages is defined. HTTP messages are then evaluated to ensure the criteria set forth by the rules are enforced.
Next-Generation Firewall	Next-generation firewalls build upon traditional firewalls to include additional protection mechanisms. Functionalities may include intrusion prevention, application firewalling, TLS/SSL-encrypted traffic inspection and more.
Cloud-Based Firewall	Firewalls that operate within the cloud on a variety of licensing terms and provide cloud-tailored features such as application control, dynamic addressing and microsegmentation. They can scale to meet the demands of the cloud.
Threat Prevention	Threat prevention services are add-on features to firewall product offerings. The services are designed to enhance firewall capabilities by adding features such as zero-day malware prevention, IDS/IPS, antivirus, DDoS protection and URL filtering. Subscription-based services can keep threat data up to date and include blacklisted IP addresses, URLs or domains.

Implementation Options in AWS

Security engineers have many options when choosing firewalls to deploy within AWS. AWS offers a native firewall solution that provides packet filtering and is integrated directly into the AWS environment. Third-party vendor solutions often offer additional features and are available from AWS Marketplace.

Customers can also engage through Consulting Partner Private Offers (CPPO) to work directly with trusted advisors to select and configure firewall solutions from AWS Marketplace. As organizations build out their cloud and cloud security strategy and plan, they may want to consider working with partners to accelerate their efforts or fill any gaps in knowledge or resources that are identified. All consulting partners may extend AWS Marketplace third-party solutions directly to customers through CPPO.²

² Consulting Partner Private Offers, <https://aws.amazon.com/marketplace/features/cpprivateoffers>

Not every organization will be able to find resources with deep cloud experience, and even experienced cloud technologists may have experience only in specific industries or with certain cloud vendors. More information on each approach is detailed in Table 2.

Table 2. Options for Choosing the Right Firewall Vendor for Use Within AWS	
Terms	Descriptions
Bring Your Own License (BYOL)	For businesses that already own firewall licenses, BYOL provides a flexible deployment option. A BYOL approach allows an organization to reassign its licenses. This approach can be ideal because the license is not tied to a specific subscription. BYOL requires that licenses be tracked. Firewalls available within AWS Marketplace may be available for use directly with AWS accounts.
Managed Firewall/ Firewall-as-a-Service	Traditionally, firewalls are a separate physical device. Managed firewalls and firewall-as-a-service offer a cloud-based rather than a device-based solution. In AWS, firewall-as-a-service offers immediate protection and, in some ways, may be more cost-effective for smaller companies that may not be able to purchase and maintain the firewall infrastructure.
Virtual Firewall Appliances	Virtual firewall appliances are installed and operate directly within the cloud. Virtual firewalls can be deployed quickly and many options are available from AWS Marketplace.
Trusted Advisors	Trusted advisors are experts in an area and can be used on a consulting basis to support selection and configuration of the optimal firewall products based on specified requirements. You can view a listing of AWS Security Competency Partners here: https://aws.amazon.com/security/partner-solutions

Needs and Capabilities: The Business Case for Firewalls and Threat Prevention in the Cloud

The perimeter is no more. But even though networks are no longer defined by their perimeters, firewall products still fill a critical role in an organization's security architecture. Firewalls have evolved from simple filtering based on IP addresses and ports. To protect today's organization, they allow security administrators to filter based on specific applications and even application functions.

Firewalls support nested policies and can be used to securely connect the data center and the cloud. Firewalls are becoming even more important as the network perimeter changes and the capabilities of attackers increase.

This section and Figure 1 detail the reasons for deploying firewalls and threat prevention services in the cloud.



Figure 1. Reasons to Deploy Firewalls and Threat Prevention Techniques in the Cloud

Blurred Line

A network perimeter is what separates the private side of a company's network from its public side. The private side is usually managed by the company, and the public network is typically managed by the provider of the network. However, with the growing popularity of mobile devices, cloud solutions and social networks, the line between private and public is increasingly blurred, making protecting the network using traditional firewalls more challenging. Mobile devices must be able to operate on networks outside the corporate firewall. Firewalls and threat prevention techniques in the cloud allow for flexibility to reconfigure according to new challenges, scalability to accommodate influxes of devices and widespread coverage beyond the physical network.

Remote Users Operate Anywhere, Anytime

Related to the disappearance of the network perimeter, more and more employees are working remotely

and accessing applications that can be hosted anywhere geographically. Traditional firewalls do not allow secure and fast connection from anywhere in the world or any time of the day. Cloud-based firewall solutions are scalable for securely tunneling all user traffic and support multifactor authentication, allowing remote users to connect via secure tunneling so that no matter where they are, their connection is secured.

Hybrid Ecosystems

As companies expand, they are turning toward hybrid ecosystems, where resources are both on premises and in the cloud. Such ecosystems reduce capital investment in physical infrastructure. Cloud-based firewalls enable hybrid ecosystems by instantiating and enforcing virtual private networks (VPNs) between the data center and the cloud. These cloud-based firewalls can be configured to scale to meet the demands of today's enterprises and can even be configured to augment the capacity of firewalls installed on premises. These cloud-based firewalls can be quickly deployed within AWS using CloudFormation templates.

Integration with SaaS Application Providers

Assuring the security of mission-critical SaaS applications can be a challenge. Cloud-based firewalls can be configured to protect against malicious attacks on these applications, and they offer features above and beyond traditional firewalls such as deep packet inspection, application-based access controls, threat prevention and zero-day malware detection.

Cost Savings

Cloud-based firewalls can be procured with flexible subscriptions. Cost models are shifting from requiring large up-front capital expenditures to monthly expenses. Cost savings can be realized through the unique licensing options available within AWS; a combination of monthly and hourly pricing supports lower-cost handling of peak demand. Additionally, when firewalls are deployed to the cloud, fewer instances may be required compared to data center installations, further reducing overall cost. Administrative costs can be lowered through automation using firewall management APIs.

Needs and Capabilities

Cloud-based firewalls provide security around the cloud implementation and support network segmentation. They enhance threat prevention capabilities.

Cloud-Based Firewalls

The need: Firewalls allow organizations to filter and log unauthorized or suspicious connections based on rules and/or behaviors. Firewalls also support network segmentation and can be used to ensure that only authorized applications or application types are run within an organization. They can also require multifactor authentication for all remote connections and can be used to detect and prevent intrusion attempts.

Capabilities

- Allow administrators to define and load policies that filter on IP addresses, ports, protocols, application types, groups and users. This capability ensures that only authorized users, communications and applications are allowed to interact with or access organizational assets, or even to limit functions within an application for some users.
- Allow administrators to segment their networks and isolate both north-south and east-west traffic. This functionality provides dynamic security across cloud/data center implementations as well as throughout the application service stack.
- Provide dynamic addressing support such as network address translation (NAT) that enables seamless integration across the cloud and data center. This support allows IP traffic across the entire ecosystem even when IP addresses change.
- Inspect encrypted traffic flowing through Transport Layer Security (TLS) tunnels. This capability mitigates the threat of an adversary passing malicious data into the network within an encrypted tunnel.
- Reduce administrative burden by providing automated policy management using well-defined APIs or providing AWS CloudFormation templates. This capability may also support touchless deployment, which significantly reduces the time needed to begin use.

Threat Prevention

The need: Threat prevention adds to a cloud-based firewall by providing advanced logging, alerting and prevention of both known and unknown threats. This feature includes services that keep firewall policy up to date with the latest threats and protects against both known and unknown malware.

Capabilities

- Provides advanced intrusion prevention capabilities that analyze, prevent and report on suspicious behavior within the system.
- Provides antivirus protections that identify and remediate malicious content based on known signatures.
- Logs events and alerts on suspicious behavior and may also support correlation across multiple firewall/threat prevention instances.
- Maintains a continually and dynamically updated threat database that includes known malware and known malicious sites and IP addresses.
- Protects the infrastructure from malware and provides advanced functionality such as DNS sinkholing.

General Cloud-Based Firewall and Threat Prevention Considerations

	Consideration	Details
	On-demand access	Today's users operate globally and 24/7. Users require secure access to their applications and data spread across the data center and the cloud.
	Hybrid ecosystems	Today's organizations use multiple infrastructures in support of their missions. Organizations spread data and applications across the data center and multiple SaaS providers. Data must be securely passed among these environments.
	Regulatory compliance mandates	Regulations mandate compliance with security and privacy requirements. Firewalls support this compliance by enforcing technical security policies that enable the confidentiality of information.
	Speed to market and agile capabilities	Organizations rely on elastic cloud services to quickly introduce new capabilities or to scale to meet demand. Cloud-based firewalls enable organizations to move quickly to meet demand and demonstrate new agile capabilities securely.
	Cost	The pay-as-you-go model enables organizations to procure cloud-based firewalls using operational dollars instead of capital expenditure (CapEx) funds. Combining hourly and annual subscriptions supports cost-effective dynamic scaling. Costs can also be saved using managed updates.
	Dynamic threat environment	Security teams are often overworked and have trouble maintaining situational awareness of the latest threats. Threat prevention services keep security teams updated on the latest in attack methods and automatically update firewall rules to guard against these new threats.

	Consideration	Details
	Application-layer support	Network communications are no longer bound to discrete service ports that can be easily filtered by a firewall. Today, most communication happens over ports 80 and 443 in the form of web traffic, leaving traditional firewalls unable to perform their functions of filtering defined IP address/port ranges. Identifying applications at Layer 7 becomes more important to safely enable the use of an application as well as reduce the attack surface.
	HTTP(S) inspection	TLS-encrypted traffic streams provide attackers with a method of gaining access to systems. Firewalls must be able to peer inside this encrypted traffic to perform filtering functions that identify the underlying application as well as any potential threats.
	Dynamic addressing	Cloud-based firewalls must be able to support environments where virtual network address ranges change on a regular basis. Dynamic addressing allows you to create policy that automatically adapts to changes—adds, moves or deletions of servers.
	Network isolation and microsegmentation	Firewalls must be able to provide network segmentation and filter traffic between trusted and untrusted environments.
	Automated policy management	Firewalls installed within the cloud must be able to be managed efficiently. APIs can support the automated management of firewall policies and enable coordination of firewall enforcement across multiple instances.
	Threat prevention	Threats change quickly, with new exploits and attack methods constantly being developed. Vendors must be able to update firewalls quickly with new information on malicious content, sites and addresses to protect the enterprise.
	Granular policy definition and enforcement	Cloud-based firewalls should be able to support policies at multiple layers of the ecosystem, including applications, application types and functions, users, networks, ports and protocols.
	Situational awareness	Firewall instances might be installed across cloud regions and within several data centers. They must be able to share logging information in standardized formats to enable situational awareness across the organization's infrastructure.

Technical Considerations

	Consideration	Details
	Single-view visibility and management	Single-view visibility makes it easier for system administrators to manage deployed firewall instances using a single management application.
	East-west traffic security	Firewalls should support the isolation of networks and security across different environments, including east-west security.
	File blocking and analysis	Threat prevention systems can block known-malicious files and analyze suspicious files before allowing them into the network. This function can keep an organization safe from the insertion of malware into the network.
	DNS monitoring	Threat prevention systems can monitor for outgoing communications to known-bad URLs and can be configured to send traffic destined to these URLs to an administrator-owned site for analysis.

Operational Considerations

	Consideration	Details
	Costs	Cloud-based firewalls can help organizations better manage their security infrastructure costs. Automated management, ease of deployment and managed updates all reduce labor costs associated with system administrators. Shifting funds from CapEx to operational budgets introduces flexibility. Combining annual subscriptions with hourly costs allows economical scalability as needed.
	Incident response	Incident response requires access to log data for situational awareness. Organizations should update incident response plans to include analysis of cloud-based firewall log information.
	Data exfiltration security	As the perimeter of the network changes and the focus shifts to data security, ensuring that data cannot be exfiltrated from the organization's network becomes critical. Threat prevention solutions flag and alert on data being sent to known-malicious sites.
	Intrusion prevention	Intrusions are blocked after evaluating traffic based on both behavior and known signatures.
	Multifactor authentication	Multifactor authentication provides an extra layer of security to VPN logins, requiring all users to use two or more forms of authentication.
	Proxy	Firewalls can act as proxies between networks, hiding the details of the private network from the outside world.

AWS Implementation Considerations

The general considerations discussed so far can help security leaders make the case for obtaining funding for the procurement of cloud-based firewalls and threat prevention services. The next section examines specific considerations for operating cloud-based firewalls within AWS. Use this section to differentiate between solutions available in AWS Marketplace.

Cloud-Based Firewalls

	Consideration	Details
	Level of AWS integration	The native AWS firewall is directly integrated with AWS services. You should ensure that AWS Marketplace firewalls have a high degree of integration with the AWS services that you use and evaluate the options for automation of deployment and update. Evaluate: • Does the firewall provide support for both virtual private cloud (VPC) and EC2 instances? • Does the firewall integrate with AWS security services such as AWS Firewall Manager, AWS Security Hub, AWS Transit Gateway and AWS GuardDuty? • Does the firewall seamlessly support high availability across multiple AWS regions? • Does the firewall offer CloudFormation templates that can reduce time to deployment?
	Policy management	Cloud-based firewalls should enable granular and automated policy management features. Evaluate: • Does the firewall support nested policies within security groups? • Does the firewall enable automated configuration of security policies? • Does the firewall support risk-based policy definitions?
	Hybrid environment support	Firewalls implement IPsec VPNs to securely network across multiple VPCs, enterprise sites and SaaS providers. Evaluate: • Does the firewall support dynamic addressing that allows you to create policy that automatically adapts to changes—adds, moves or deletions of servers? • Does the firewall support networking across multiple VPCs?
	Logging	Logs provide a vital resource for incident response and forensics. All firewalls should provide logging features. Evaluate: • Does the firewall offer a solution that allows for aggregation of logs across multiple firewall instances? • Does the firewall integrate with AWS logging mechanisms?
	AWS security competency approval	AWS security competencies for infrastructure security products provide a degree of confidence that the firewall meets minimum security standards for operation within AWS. Evaluate: • Does the firewall have AWS security competency approval? • Does the firewall meet other security standards and best practices?
	Application control	Firewalls should provide administrators with the capability to set policy based on the organization's specific needs. This capability includes filtering on approved applications and nesting policy within security groups. Evaluate: • Does the firewall support filtering based on app ID to permit only approved applications within the network? • Does the firewall support dynamic application filters and application groups that restrict the types of applications authorized on the network? • Does the firewall support dynamic profiling, allowing the firewall to learn the typical behavior of the application over time?

Cloud-Based Firewalls (continued)

	Consideration	Details
	Separation of trusted and untrusted zones	Firewalls must be able to segregate both north-south and east-west traffic. This segregation allows untrusted zones (such as development) to interact with trusted zones (such as production), and supports processes such as DevOps. Evaluate: • Does the firewall filter across trusted and untrusted zones? • Does the firewall support micro-segmentation and isolation of subnetworks?
	Management of multiple firewall instances	Many firewall vendors provide software that allows for the seamless management of multiple firewall instances. Evaluate: • Does the firewall include software that can manage all of the firewall instances in the cloud? • Does the firewall management software allow you to push policies and perform updates to device configurations?
	Scalability	Cloud-based firewalls should support elastic expansion, allowing them to scale automatically to meet the demands of users. Evaluate: • Does the firewall scale automatically? • Can you use the firewall to augment data center installations to support peak demand (e.g., cloudbursting)?
	Dynamic reporting	Reporting provides administrators with insight into trends as events occur across the network. Cloud-based firewalls should provide insightful reporting features. Evaluate: • Does the firewall provide reporting that allows for analysis of incoming requests? • Does the firewall provide reporting that tracking of trends in violations?

The above considerations are based on integration of firewall capabilities within an AWS environment. Organizations may not need all of the capabilities discussed here, but they can review these considerations and determine what is needed based on their specific requirements. A critical consideration, however, is the capability to seamlessly integrate with AWS services. Any solution selected from AWS Marketplace should provide this baseline capability.

Threat Prevention

Threat prevention is critical to keep organizations ahead of the dynamically changing threat landscape. Threat prevention techniques incorporate the latest threat intelligence data and dynamically update policies to guard against the latest attack methods and malicious sites. Threat prevention services can provide file-blocking features, keep data from leaving the network, and identify and prevent intrusions.

Threat Prevention

	Consideration	Details
	Cloud context support	Threat prevention is based heavily on the ability to acquire relevant information on the latest threats, threat actors and their capabilities. Ensure that the threat prevention services you procure within AWS are supported by top-quality threat intelligence feeds. Evaluate: • Is the threat intelligence data timely? • Is the threat intelligence data relevant to your organization's mission? Threat prevention services should keep customers up to date on the latest threats to their systems.
	Performance and efficiency	Evaluate: • Does the threat prevention service provide a listing of known-bad addresses and sites? • Does the threat prevention service automatically update new malware signatures? • Does the threat prevention service automatically update firewall rules based on known malicious activity? • Does the threat prevention service have the ability to perform DNS sinkholing or DNS security?
	Deployment	Firewalls incorporating threat prevention should be capable of creating a baseline of behavior and alerting on anomalies. Evaluate: • Does the threat prevention service analyze logs, correlate events and block/alert on suspicious activity? • Does the threat prevention service support behavioral analysis? • Does the threat prevention service scan all traffic, including applications, users and content?
		Threat prevention services should incorporate antivirus support that includes maintaining an updated list of signatures. Evaluate: • Does the threat prevention service incorporate network antivirus features? • Does the threat prevention service provide file-blocking and analysis capabilities?
		Threat prevention services should provide features that keep data from leaving the network. Evaluate: • Does the threat prevention service support DNS monitoring and redirection to an administrator-specified site? • Does the threat prevention service flag on traffic destined to known malicious domains?

The above should be taken into consideration when choosing threat prevention services to add on to your firewall platform procurement within AWS.

Making the Choice

A simple Analysis of Alternatives (AoA) will allow your organization to objectively compare the products available in AWS Marketplace against one another and against the native AWS firewall service. An AoA consists of multiple steps that include:

1. Review this guide and identify your organization's specific requirements.
2. Weigh the requirements according to the importance to your organization. For example, weigh critical requirements as "high" and desired requirements as "low." Cost should also be considered as a factor in the evaluation.
3. Review the capabilities of the native AWS firewall.
4. Compile a list of vendor firewall/threat prevention offerings from AWS Marketplace.
5. Evaluate each firewall/threat prevention offering against selected requirements.
6. Score each of the products against each requirement.
7. Calculate the sum score for each offering and select the product with the highest score.

Organizations can also opt to contract through AWS Marketplace CPPO to perform this analysis of alternatives. Choosing this approach is often optimal based on the level of expertise available through these partner organizations.

Conclusion

Options for cloud-based firewalls for use in an AWS deployment include native AWS offerings and third-party products offered in AWS Marketplace. An analysis of the available options based on the considerations in this paper will allow for the selection of a firewall that meets the unique requirements of any organization. Critical considerations when choosing firewall and threat prevention capabilities include the abilities to separate trusted and untrusted zones, evaluate encrypted traffic, perform behavioral analysis, operate across hybrid environments and integrate directly with AWS services. To perform this analysis, identify firewall and threat prevention options available today in AWS Marketplace and evaluate each against the criteria in this paper.

Performing a formal analysis of alternatives will support an objective determination of the best technology solution. Alternatively, organizations can reach out to trusted third-party Consulting Partners to customize a firewall and threat prevention approach for security within the cloud. Visit the AWS Security Competency Partners page³ for more information.

³ AWS Security Competency Partners, <https://aws.amazon.com/security/partner-solutions>

About the Author

Brian Russell is the Chair of the Cloud Security Alliance (CSA) Internet of Things (IoT) Working Group and founder at TrustThink, LLC where he leads security engineering for autonomous vehicles and smart devices. He was previously Chief Engineer for Cyber Security Solutions at Leidos - a Fortune 500 Government Contractor. In that role he led Research and Development (R&D) for secure cloud systems, permissioned blockchain networks, and cryptographic key management. Brian is an adjunct professor with the University of San Diego (USD) in the graduate Cyber Security Operations and Leadership Program and co-author of the book Practical Internet of Things Security.

Chapter 22: Solution Guidance for Endpoint Security in AWS



David Hazar

SANS Instructor

"Endpoint security in the cloud is an emerging challenge that will continue to evolve as more and more cloud customers choose to update their architecture and design to take advantage of nontraditional or non-infrastructure as a service (IaaS) cloud services and as vendors and cloud providers create additional cloud-native or cloud-optimized solutions.

Organizations will need to continuously evaluate these capabilities as they update their cloud architecture and design, paying particular attention to the capability or product's use of cloud context, efficiency, and ease of use and integration. Direct and indirect costs will also need to be considered and measured to ensure these costs can be supported as cloud usage increases and that the total cost of endpoint security is known."

Introduction

Endpoint security options and products are continuing to mature. Enterprises and other organizations are moving away from point solutions—antivirus (AV) or anti-malware, host-based intrusion detection systems (HIDSs), file integrity monitoring (FIM) and application whitelisting—toward more robust endpoint protection platforms (EPPs).

And many of those EPPs include new, advanced endpoint detection and response (EDR) capabilities. This move is similar to other efforts to consolidate the functionality of multiple security capabilities into a single solution or platform to make it easier for organizations to implement and maintain these technologies.

Just as these firewalls bring the capabilities of many different security appliances into a single solution, EPPs bring the capabilities of many endpoint security agents into a single agent, or at least a single management platform.

Gartner describes EPPs as “a solution deployed on endpoint devices to prevent file-based malware attacks, detect malicious activity and provide the investigation and remediation capabilities needed to respond to dynamic security incidents and alerts.” A wide range of products and solutions falls into this category, in part because there is no strict definition of required capabilities for them to be considered EPPs. That’s why you will find many traditional point solutions from recognizable vendors included in this category, albeit bundled together with some new solutions or with the addition of some new capabilities. You will also find more recent entrants into the endpoint security market that may have new, innovative approaches to endpoint security but may also lack maturity in more traditional detection and response capabilities.

“Just as next-generation firewalls bring the capabilities of many different security appliances into a single solution, EPPs bring the capabilities of many endpoint security agents into a single agent.”

Selecting and implementing endpoint security in hybrid architectures can be a time consuming and confusing process. In this paper, we present what customers should consider when evaluating endpoint security technology in the cloud. We discuss a high-level strategy for evaluating these solutions and then discuss implementation options that organizations need to consider when planning to implement these technologies in Amazon Web Services (AWS). We also review why businesses may choose to implement endpoint security in the cloud along with the various needs and capabilities associated with different endpoint security solutions. Lastly, we discuss some of the considerations that should be part of the evaluation process for endpoint security in general, but then take a closer look at the considerations specific to implementing endpoint security in AWS.

Not all companies may choose or be able to implement endpoint security for all of their cloud workloads. Because much of the technology associated with endpoint security is installed and runs as an agent, infrastructure-as-a-service (IaaS) cloud workloads are the most obvious candidates. In AWS, endpoint security solutions typically work with EC2 Instances or virtual machines (VMs) created on VMware Cloud on AWS. While these technologies could technically also be leveraged within containerized environments, such a situation is less typical and other container security technologies may be better suited in this type of environment. This paper focuses on implementation via instances or VMs, but most of the considerations still apply to a containerized environment.

Some cloud service types, such as platform-as-a-service (PaaS), function-as-a-service (FaaS) and software-as-a-service (SaaS), are not supported by many endpoint security technologies. However, the considerations outlined in this paper can help customers determine what protections vendors provide for these service types. There is also a case to be made for leveraging the cloud shared-responsibility model to reduce an organization's security burden if the risk for those workloads does not merit the increased visibility or if an organization feels it cannot provide better protection than the cloud vendor even with the increased visibility. In these situations, leveraging PaaS, FaaS and SaaS cloud services can help.

Understanding Your Needs

In order to evaluate endpoint security, organizations need to have a solid understanding of what capabilities are must-haves versus nice-to-haves to provide the level of protection and visibility they desire. They must also consider how the endpoint security program will be implemented, operated and maintained. Organizations should avoid purchasing technology if there is not sufficient support, funding, resourcing and processes in place to successfully implement, operate and maintain the technology for years.

After the organization determines and ranks capabilities, it needs to look at existing endpoint security technology, people and processes to understand what is currently in place and whether it is well suited for the cloud. Then, it should investigate alternative technologies, including any cloud-optimized solutions, and catalog the resources and skills that will be required, along with the policies, standards and processes that may need to be updated. This investigation will not be a one-time exercise; these points will be revisited many times throughout the evaluation process before making a choice.

“Organizations should avoid purchasing technology if there is not sufficient support, funding, resourcing and processes in place to successfully implement, operate and maintain the technology for years.”

Implementation Options in AWS

When the cloud was new, the only real option was to leverage technology similar to what an organization was already using on premises, if not the same technology. If you already have a successful and functional on-premises program, this can be an attractive option, but it is not the only option. Review the different options you have available, including cloud-optimized, managed services and licensing options. Then, once you have a rough idea of how you would like to implement endpoint security in AWS, it is time to start building a business case.

Cloud-Optimized

Organizations may want to look at cloud-optimized solutions for endpoint security in the cloud. Traditional endpoint security technology is typically not performance-friendly. In on-premises environments, the costs for this overhead are not always as easy to see or calculate because there is usually excess capacity that can be used to compensate for the overhead. In the cloud, however, with on-demand pricing and the detailed metrics, the cost of this overhead is much easier to understand. Many cloud-optimized endpoint security tools focus on creating lightweight agents that offload the processing of data and events to other resources or even to a separate, vendor-maintained cloud environment.

Managed Services

Another option for implementing endpoint security in the cloud is to leverage a managed service provider that has experience implementing and maintaining these solutions in the cloud. Using such a provider can be a promising option for many organizations but is especially attractive for organizations that have limited cloud experience or that do not already have endpoint security capabilities. Another advantage of managed service providers is that they typically provide skilled resources and bring with them proven processes and existing cloud vendor contacts and relationships to accelerate implementation and add value quickly. They may also supplement the endpoint security technology with human-assisted analysis, custom development or configuration, and even incident response capabilities. These managed service providers can even extend AWS Marketplace solutions directly to customers through Consulting Partner Private Offers and assist with evaluating licensing options.

Licensing Options

When considering how to implement endpoint security in the cloud, also consider how to license any chosen technology. If you are planning on using existing on-premises endpoint security capabilities, your organization may already have favorable licensing and it may make sense to follow a bring-your-own-license (BYOL) model.

Maybe endpoint security is new to your organization, or maybe you want to evaluate a technology without implementing it more broadly. Perhaps you determine you need a different technology for the cloud or your organization favors on-demand pricing or operational cost structures. If any of those scenarios apply to you, you'll be relieved to learn that many of the products are available with on-demand pricing from AWS Marketplace. (AWS Marketplace can still be leveraged for many of these technologies following the BYOL approach as well).

Needs and Capabilities

Cloud architecture differs from what we are used to in our on-premises environments. In the cloud, almost everything is software-defined—and we do not have complete visibility into our resources and surrounding infrastructure. Also, because commissioning and decommissioning resources are so easy to do and costs are typically accrued based on the amount of time the resource is running, cloud resources tend to have much shorter lifecycles. The capabilities surrounding forensics in the cloud are also much less mature than for on-premises environments, and leveraging endpoint security can provide valuable threat intelligence for an organization's cloud ecosystem that it may not be getting from its PaaS, FaaS and SaaS workloads.

Next, we look at some of the solutions or capabilities that may exist within endpoint protection platforms and then move on to the topics organizations should consider when preparing to implement endpoint security in the cloud.

Needs and Capabilities

Note the overlap between the solutions listed below. For example, EDR solutions may provide many of the same capabilities as AV/anti-malware or HIDS solutions. Some AV solutions may also include behavior monitoring, and both HIDS and EDR solutions will most likely perform FIM. This overlap in capabilities will be one of the considerations for organizations that choose to utilize more than one solution.

Endpoint Detection and Response

The need: Identifying and protecting against unknown threats

Capabilities

- Detecting security incidents
 - Behavior monitoring
 - Analytics
 - Sandboxing
- Containing the incident at the endpoint
- Investigating security incidents
- Providing remediation guidance

Antivirus/Anti-malware

The need: Identifying and protecting against known threats

Capabilities

- Detecting viruses and malware

- Signature analysis
- Behavior monitoring
- Blocking and quarantining the virus or malware
- Alerting users and administrators of infection

Host-based Intrusion Detection

The need: Identifying indicators of compromise

Capabilities

- Detecting suspect behavior
 - Behavior monitoring
 - Traffic analysis
- FIM
- Alerting users and administrators of suspect behavior

File Integrity Monitoring

The need: Identifying changes to critical or sensitive files

Capabilities

- Collecting and storing signature data for policy-defined files
- Offering interval-based or real-time signature validation
- Alerting users or administrators when tracked files are modified

Application Whitelisting

The need: Only allowing approved or authorized, signed software to execute

Capabilities

- Authorizing software or software signing certificates via policy
- Applying policies to resources
- Blocking or alerting when unauthorized software executes

General Cloud Endpoint Security Considerations

Regardless of the endpoint security technology or cloud vendor selected, some general business, technical and operational considerations are associated with implementing endpoint security in the cloud. The following sections highlight many of these considerations.

Business Considerations

	Consideration	Details
	Policies and standards	Traditional endpoint security requirements in policies and standards may not be achievable in the cloud, may not function as intended or may not be cost-effective. Organizations will need to evaluate cloud capabilities to determine what changes need to be made to ensure that compliance with policies and standards is achievable.
	Governance model	Every organization has a unique governance model. Some organizations have very centralized governance over endpoint security, whereas others may follow a more decentralized approach. Organizations will need to decide whether to centralize or decentralize governance over cloud endpoint security. Then, they must determine whether existing governance models used for traditional endpoint security can be extended to the cloud or whether a cloud-specific model is required. Consider that cloud workloads can easily span the globe and that data residency and visibility restrictions may apply in certain regions.
	Reporting and metrics	Providing the right metrics, key performance indicators (KPIs) and key risk indicators (KRIs) to the right stakeholders may require changes that account for cloud architecture. Organizations will need to define reporting requirements specific to cloud workloads and evaluate products against these requirements
	Funding and support	Funding and support for cloud endpoint security may not currently be available. Organizations may not understand the shared responsibility model as it pertains to cloud usage and may assume that endpoint security is provided by the cloud vendor. Organizations will need to understand the requirements and determine the appropriate funding and support model. What is required may differ based on the implementation model the organization chooses (for instance, traditional vs. cloud, BYOL vs. on-demand).
	Risk classification	Not all workloads share the same risk profile. It is important that organizations consider the risk associated with different cloud workloads to enable them to implement controls based on risk. If cost is not a consideration or if the risks are similar for all workloads, then a single approach to endpoint security may be appropriate. If risks vary greatly among workloads or costs are high, however, an organization will need to understand the various risk profiles to determine where to focus or require endpoint security or what to require for each profile.

Technical Considerations

	Consideration	Details
	Endpoint security capabilities	As organizations update policies and standards to address cloud workloads, they should also identify the technologies they need to comply with these new requirements. Some organizations may choose to be prescriptive about the technologies they use, whereas others may define the required capabilities and allow individual cloud operations teams to select their own technologies as long as they can validate compliance with requirements.
	Supported technology	Some technologies may not be supported for all cloud services or for all platforms running on cloud services. Organizations will need to decide whether they will allow the use of services and platforms that do not support endpoint security requirements, and if so, under what conditions. These decisions should be documented and maintained so they may be consistently applied throughout the organization.
	Agent-based technologies	No matter how lightweight, agent-based technologies decrease performance (most cloud endpoint security technologies are agent-based). In the cloud, they increase costs. Organizations may have a restriction on the number of agents that can be installed on each cloud resource. Organizations need to determine how many non-endpoint security agents are already in place to decide whether they need to consider an increase in their limits. They may also have a specific overhead allowance for agents, which needs to be evaluated during any proof of concept. Performance should be assessed before purchase, before upgrades, when configurations change and at regular intervals. Metrics should include overhead and performance monitoring.
	Active vs. interval-based or asynchronous detection and response	Technologies that provide active detection and response may require more overhead than technologies that scan at given intervals, during off-peak hours or asynchronously via out-of-band analysis engines. Organizations need to decide whether active detection and response are required or acceptable based on their cloud architecture. In particular, the longevity of cloud resources may affect this decision. Short-lived cloud resources may require more active defenses.
	Secure communication	Endpoint security solutions all typically communicate with external components or services. The external services could provide product updates or configuration data. They may also be involved in the analysis of data from the target system. Organizations need to ensure that external communication is authenticated and secured.

Operational Considerations

	Consideration	Details
	Operational responsibility and model	Operation of cloud resources is substantially different from traditional infrastructure operations. This difference may help determine who is responsible for implementing and configuring endpoint security capabilities. Organizations need to decide how best to implement and configure endpoint security technology and determine which group(s) should be responsible for this task. They also need to determine whether operations should be centralized or decentralized.
	Monitoring and response	While implementation and configuration of endpoint security capabilities may be assigned to an existing cloud operations team, monitoring may be the responsibility of others. Response could be the responsibility of either team. Organizations need to determine who will be responsible for monitoring and responding to endpoint security events. They will also need to evaluate what orchestration and automation of technology, people and processes can be leveraged or integrated into the final solution.
	Processes and procedures	Organizations may have very specific processes and procedures for dealing with endpoint security events related to their traditional on-premises infrastructure. It is likely, however, that these processes and procedures will be different in the cloud. Organizations need to create new operational processes and procedures for endpoint security in the cloud, considering any changes they have made to policies and standards related to cloud or endpoint security in the cloud.

AWS Implementation Considerations

The general considerations discussed so far can help organizations lay the groundwork as well as secure funding and support for cloud endpoint detection. Now let's take a more detailed look at some specific considerations an organization will need to evaluate before implementing these solutions in AWS.

Endpoint Detection and Response

The advantage of EDR solutions is that they focus on adding capabilities that allow them to identify unknown threats. If your organization's threat profile includes targeted attacks or advanced threat actors, consider endpoint protection platforms that excel in EDR. You may also consider EDR for high-risk workloads or for performance reasons, because many of these solutions offload processing to other resources. False positive rates may be higher for EDR than some other types of tools.

	Consideration	Details
	Cloud context support	Due to the dynamic nature of the cloud, a resource that existed a few hours ago may not exist now. Because many EDR technologies perform analysis of data or binaries external to the resource itself, there is a chance that when analysis is completed the resource may no longer exist. Evaluate: • The additional cloud context (specifically, tags or image IDs) that is captured, retained and used by EDR technology to allow correlation of findings and behavior with resources and the images and image versions used to create those resources • The special concerns associated with studying resources that have potentially replaced the original resource from which data was gathered
	Performance and efficiency	Many EDR platforms claim to have lightweight agents that offload analysis and processing tasks to other systems. Customers should analyze the impact and performance on production workloads. Due to the offload architecture, these technologies typically send data and binaries to separate systems or to the vendor's cloud infrastructure to perform analysis. Depending on the cloud regions in use, the transfer of data and binaries to external resources could affect both the performance and the cost of the technology. In addition, depending on the architecture, analyzing the same data and binaries from multiple systems may add additional processing time and reduce efficiency. Evaluate: • The architecture of the tools under consideration • Performance (CPU, memory, storage and bandwidth utilization) when used with production workloads • The amount of data and binaries that will be transferred and to what location(s) the data is being transferred • Potential impacts on cost and performance due to bandwidth • Performance impact of latency between all cloud regions in use and any identified external resources • Efficiency of coordination between agent and analysis engine(s) and efficiency of threat data distribution • Support for data compression

Endpoint Detection and Response (continued)

	Consideration	Details
	Deployment	EDR platforms may require the implementation of multiple components, and these components may need to be installed in multiple zones or regions to support distributed cloud environments. They also require the implementation of agents on the supported endpoints. Evaluate: • The architecture of the tools under consideration • The installation and configuration procedures for each component and agent • The availability of managed or SaaS components or preconfigured appliances from AWS Marketplace • Effectiveness and responsiveness of support • Any vendor requirements for the use of professional services for installation or configuration • Integration with other AWS technologies for deployment or validation of agent deployment (AWS Systems Manager, AWS Config, Amazon CloudWatch)²
	Configuration and maintenance	In order to improve the quality of detection and response, EDR technologies may require extensive configuration and maintenance. Customization of detection rules and response scripts may be available depending on product. In addition, EDR components and agents will need to be upgraded and may also require updates to datasets used for analysis. Evaluate: • The architecture of the tools under consideration • The upgrade procedures • The procedures for updating any datasets leveraged by analysis engines or agents • Reporting, metrics or alerting available for any out-of-date components, agents or data • Communication protocols and paths to understand required firewall and ACL changes along with any VPC peering or cross-account access • Any vendor requirements for the use of professional services for upgrades or updates • Accessibility to detection rules, scripts and other configuration details (open or proprietary) • Whether the platform allows customers to build or create their own rules • Level of effort to perform customizations to rules, scripts or configurations or to create new rules • Integrations with other AWS technologies (such as AWS Config, AWS Lambda) or configuration management tools (Puppet, Chef, Ansible, SaltStack, CFEngine) to perform updates or upgrades or apply configurations • Secure configuration guides and best practices
	Detection	Because EDR technologies support detecting both known and unknown threats, organizations should evaluate their effectiveness as part of the selection process. Evaluate: • Detection rate of any well-known or unknown malware samples, if your organization practices malware analysis and has appropriate analysis environments • Detection methods employed • Available benchmarks or comparisons by third-party evaluators • Product reviews and customer forums • Customer references • Whether detection is real-time, interval-based, asynchronous or configurable for each detection mechanism supported • Whether detection includes detection of non-file-based malware (such as memory resident malware)

This paper mentions product names to provide real-life examples of how visibility tools can be used. The use of these examples is not an endorsement of any product.

Endpoint Detection and Response (continued)

	Consideration	Details
	Integration	Many EDR platforms also integrate with other business and security platforms. Understand what integrations are supported out-of-the-box and the level of effort required to build custom integrations. Evaluate: • Supported plugins and integrations with business and security platforms in use by the organization (such as AWS, ticketing, SIEM, incident response, threat intelligence) and the capabilities of these plugins and integrations • API support (such as API-first, REST API available, programmatic API available) • Whether the platform allows the customer to build custom plugins or integrations • Level of effort required and technology (languages, frameworks, and the like) supported when building custom plugins or integrations
	Reporting, metrics and alerting	EDR platforms have response capabilities, but not all rules trigger a response. Accessing and viewing what these platforms detect and the actions they take is critical to the security of the organization's endpoints. Taking that action can also aid in the identification of rules or configurations that require modification and can also assist in troubleshooting production incidents that may be caused by the EDR platform (false positive detection and response). Evaluate: • Support for centralized logging technologies and communication protocols, including integration with any existing or proposed SIEM technology • Out-of-the-box reports and dashboards against current program requirements • Ability and level of effort required to create custom measures and metrics • Alerting mechanisms and ability to create or modify alerts • Supported reporting and alerting formats and delivery mechanisms • Integration with AWS reporting and alerting tools (such as AWS Security Hub, Amazon CloudWatch Events, Amazon Simple Notification Service [SNS]) • Support for data aggregation across regions • Supported data export formats
	Response capabilities	Another distinguishing factor when evaluating EDR technologies is what response capabilities are available in the platform. Understand not only what response abilities exist for both human-assisted and automated response but also what expertise is required to set up, configure and maintain these capabilities. Evaluate: • Out-of-the-box response capabilities and features • Technologies and languages supported for automated response • Organization's ability to support automation through identified technologies and languages • Auditing and tracking of response actions • Integration with AWS response capabilities and APIs

EDR platforms are becoming more popular as organizations strive to protect themselves against emerging threats and want to acquire more active response capabilities. We have also seen, however, that companies utilizing these technologies are still susceptible to security breaches. Implementing an EDR platform requires more than just the licensing and implementation of the technology components. It requires active monitoring, response, reconfiguration and maintenance. Make sure your organization is aware of the true costs of ownership: training requirements, resource requirements, integration

requirements and the costs to update policy, standards, processes and procedures. Also, make sure to thoroughly evaluate reporting, monitoring and alerting capabilities, because these are the most likely to require customization or integration work.

Antivirus/Anti-malware

The advantage of AV solutions is that they are typically mature products that excel at identifying known viruses and malware using signature-based and other techniques. Although attackers can easily evade these detections, positive identification from these tools indicates a real threat, and false positive rates are low when organizations use signature-based detection. Consider mature AV products if EDR technologies are prohibitively expensive, incapable of detecting known threats, difficult to tune or drags on performance. You still need to complete a performance analysis for AV capabilities because the resource requirements will vary based on architecture and supported detection mechanisms.

	Consideration	Details
	Cloud context support	If investigation of AV detections is delayed, the resource(s) affected may no longer exist in your cloud environment. Also, if a virus or worm is spreading throughout your environment, understanding more about the cloud asset can help speed response to the threat. Evaluate: • The additional cloud context (such as tags or image IDs) that is captured and retained by AV technology to allow correlation of detections with resources and the images and image versions used to create those resources • The special concerns associated with studying resources that have potentially replaced the original resource from which data was gathered
	Performance and efficiency	Traditional AV agents are not known for being lightweight and are much more likely to store and process data on the cloud resource itself. Consider how this will affect instance sizing and storage requirements. Evaluate: • The architecture of the tools under consideration • Performance (CPU, memory, storage and bandwidth utilization) when used with production workloads • Amount of data sent and received from management console(s) • Amount of data stored on disk (such as signature database, logs, quarantine) • Support for data compression
	Deployment	AV software requires agents and may also report data back to a management console. Update servers may also be used to distribute updates to the software and signature database. Evaluate: • The architecture of the tools under consideration • The installation and configuration procedures for agents and any management infrastructure • The availability of managed or SaaS components or preconfigured appliances from AWS Marketplace • Effectiveness and responsiveness of support • Any vendor requirements for the use of professional services for installation or configuration • Integration with other AWS technologies (such as AWS Systems Manager, AWS Config, Amazon CloudWatch) for deployment or validation of agent deployment

Antivirus/Anti-malware (continued)

	Consideration	Details
	Configuration and maintenance	Traditional AV products are not as configurable as EDR platforms, but it is still important to understand and review configurations on a regular basis. Reviews should be required on changes to the default configuration. Evaluate: • The architecture of the tools under consideration and the upgrade/update procedures • The procedures for updating signatures • Communication protocols and paths to understand required firewall and ACL changes along with any VPC peering or cross-account access • Reporting, metrics or alerting available for any out-of-date agents or signatures • Any vendor requirements for the use of professional services for upgrades or updates (not common) • Ability to customize scan intervals or manage exclusions • Whether the platform allows customers to add their own signatures • Availability and content of secure configuration guides and best practices
	Detection	Traditional AV technologies focus primarily on known threats. It is important for organizations to evaluate their effectiveness as part of the selection process. Evaluate: • Detection methods of any known malware samples available, if your organization practices malware analysis and has appropriate analysis environments • Available benchmarks or comparisons by third-party evaluators • Product reviews and customer forums • Customer references • Whether detection is real-time, interval-based, asynchronous or configurable for each detection mechanism supported • Whether detection includes detection of non-file-based malware (memory resident malware, for example)
	Integration	Traditional AV platforms have historically operated independently of other technologies and systems with the exception perhaps of log aggregation technologies, but it is still important to understand what integrations are supported out-of-the-box and the level of effort required to build any custom integrations. Evaluate: • Supported plugins and integrations with business and security platforms in use by the organization (such as AWS, ticketing, SIEM, incident response, threat intelligence) and the capabilities of these plugins and integrations • API support (API-first, REST API available, programmatic API available, to name a few) • Whether the platform allows the customer to build custom plugins or integrations • Level of effort required and technology (such as languages or frameworks) supported when building custom plugins or integrations

Antivirus/Anti-malware (continued)

	Consideration	Details
	Reporting, metrics and alerting	AV software will detect and attempt to neutralize a high percentage of well-known threats in your environment. Nevertheless, implement adequate reporting, metrics and alerting to respond quickly when you see new threats in your environment, because the extent of the automated response may be limited to killing processes and quarantining malware. Enhance the effectiveness of the technology by supporting defined standards and goals. Evaluate: • Out-of-the-box reports and dashboards against current program requirements • Ability and level of effort required to create custom measures and metrics • Alerting mechanisms and ability to create or modify alerts • Supported reporting and alerting formats and delivery mechanisms • Integration with AWS reporting and alerting tools (such as AWS Security Hub, Amazon CloudWatch Events, Amazon SNS) • Support for data aggregation across regions • Supported data export formats

There are many mature options when evaluating AV solutions. However, not all of these vendors have focused on optimizing their technology for the cloud. Take this into consideration when evaluating your current on-premises AV technology against other options for implementation in cloud environments. Performance and reporting may be the biggest considerations when implementing AV in the cloud.

Host-based Intrusion Detection

HIDS capabilities will be included with EDR solutions and bundled with other EPPs even if they may not advertise themselves as EDR solutions. Consider an EPP or product that focuses on HIDS if 1) EDR is prohibitively expensive or negatively affects performance, and 2) you want to detect indicators of compromise (IoCs). You will still need performance analysis for HIDS capabilities, because the resource requirements will vary based on the detection mechanisms supported and your architecture. HIDS may also offer more visibility into network traffic, depending on product capabilities.

Host-based Intrusion Detection

	Consideration	Details
	Cloud context support	If investigation of HIDS detections is delayed, the resource(s) affected may no longer exist in your cloud environment. Understanding more about the cloud asset can help speed response to the threat. Evaluate: • The additional cloud context (tags or image IDs) that is captured and retained by HIDS technology to allow correlation of detections with resources and the images and image versions used to create those resources • The special concerns associated with studying resources that have potentially replaced the original resource from which data was gathered
	Performance and efficiency	Because of the move to EDR, traditional HIDS agents may not be optimized for cloud. Consider how this will affect instance sizing and storage requirements. Evaluate: • The architecture of the tools under consideration • Performance (CPU, memory, storage and bandwidth utilization) when used with production workloads • Amount of data sent and received from management console(s) • Amount of data stored on disk (logs) • Support for data compression

Host-based Intrusion Detection (continued)

	Consideration	Details
	Deployment	HIDS software requires agents to identify indicators of compromise and may also report data back to a management console. Evaluate: • The architecture of the tools under consideration • The installation and initial configuration procedures for agents and any management infrastructure • The availability of managed or SaaS components or preconfigured appliances from AWS Marketplace • Effectiveness and responsiveness of support • Any vendor requirements for the use of professional services for installation or configuration • Integration with other AWS technologies (such as AWS Systems Manager, AWS Config, Amazon CloudWatch) for deployment or validation of agent deployment
	Configuration and maintenance	HIDS agents and corresponding policies or rules will need to be tuned to eliminate false positives and may require custom rules or policies to monitor specific configurations or logs. They will also require regular maintenance and updates/upgrades. Evaluate: • The architecture of the tools under consideration • The upgrade procedures • The procedures for updating any datasets leveraged by agents • Reporting, metrics or alerting available for any out-of-date agents or policies • Communication protocols and paths to understand required firewall and ACL changes along with any VPC peering or cross-account access • Any vendor requirements for the use of professional services for upgrades or updates • Accessibility to detection rules, scripts and other configuration details (open or proprietary) • Whether the platform allows customer to build or create their own rules • Level of effort to perform customizations to rules, scripts or configurations or create new rules • Integrations with other AWS technologies (such as AWS Config and AWS Lambda) or configuration management tools (Puppet, Chef, Ansible, SaltStack, CFEngine) to perform updates or upgrades or apply configurations • Secure configuration guides and best practices
	Detection	HIDS technologies require the implementation of agents on the supported endpoints. They may also require the provisioning and deployment of management consoles and centralized update servers or appliances. Evaluate: • Detection methods employed • Data and services included for monitoring and detection • Available benchmarks or comparisons by third-party evaluators • Product reviews and customer forums • Customer references
	Integration	HIDS software may support integration with other reporting and alerting capabilities. Evaluate: • Supported plugins and integrations with business and security platforms in use by the organization (such as AWS, ticketing, SIEM, incident response, threat intelligence) and the capabilities of these plugins and integrations • API support (such as API-first, REST API available, programmatic API available) • Whether the platform allows the customer to build custom plugins or integrations • Level of effort required and technology (languages and frameworks) supported when building custom plugins or integrations

Host-based Intrusion Detection (continued)

	Consideration	Details
	Reporting, metrics and alerting	HIDS technologies focus on detection. For that reason, reporting, alerting, monitoring and response procedures are crucial. Evaluate: • Out-of-the-box reports and dashboards against current program requirements • Ability and level of effort required to create custom measures and metrics • Alerting mechanisms and ability to create or modify alerts • Supported reporting and alerting formats and delivery mechanisms • Integration with AWS reporting and alerting tools (such as AWS Security Hub, Amazon CloudWatch Events, Amazon SNS) • Resources and processes to support monitoring of reports and response to alerts • Support for data aggregation across regions • Supported data export formats

HIDS can provide insight into what is happening on your endpoints when more advanced endpoint detection and response is not available. However, if cloud endpoints have short lifecycles, HIDS may not provide as much value unless enough cloud context is available to determine which detections or events are relevant to similar cloud endpoints or the cloud endpoint that replaced the endpoint on which the initial event occurred.

File Integrity Monitoring

FIM may be included in many of the other EPP solutions, but you may consider it as a point solution if integrity is significantly more important than confidentiality and availability, and if the capabilities of the solutions included in your EPP do not meet your needs. FIM may become less important as organizations move toward more immutable workloads, where most sensitive files reside on read-only portions of the file system and more consistently leverage PaaS for back-end storage technologies (such as Amazon RDS, Amazon S3).

	Consideration	Details
	Reporting, metrics and alerting	File integrity monitoring typically affects performance much less than other endpoint security technologies because it is focused only on the integrity of files. Performance should still be evaluated before using these technologies in the cloud, especially when other security agents are also installed. Evaluate: • The architecture of the tools under consideration • Performance (including CPU, memory, storage and bandwidth utilization) when used with production workloads • Amount of data sent and received from management console(s) • Amount of data (such as file hash/signature database, logs) stored on disk • Support for data compression

File Integrity Monitoring (continued)

	Consideration	Details
	Deployment	FIM software requires agents to identify changes to monitored files and may also report data back to a management console. Evaluate: • The architecture of the tools under consideration • The installation and initial configuration procedures for agents and any management infrastructure • The availability of managed or SaaS components or preconfigured appliances from AWS Marketplace • Effectiveness and responsiveness of support • Any vendor requirements for the use of professional services for installation or configuration • Integration with other AWS technologies (such as AWS Systems Manager, AWS Config, Amazon CloudWatch) for deployment or validation of agent deployment
	Configuration and maintenance	Configuration and maintenance of FIM software may be less cumbersome than the other solutions we have discussed, but all solutions require some degree of configuration and maintenance. Evaluate: • The architecture of the tools under consideration • The upgrade procedures • Reporting, metrics or alerting available for any out-of-date agents • Communication protocols and paths to understand required firewall and ACL changes, along with any VPC peering or cross-account access • Any vendor requirements for the use of professional services for upgrades or updates • Level of effort to configure policy that determines which files to monitor • Integrations with other AWS technologies (such as AWS Config and AWS Lambda) or configuration management tools (Puppet, Chef, Ansible, SaltStack, CFEngine) to perform updates or upgrades or apply configurations • Secure configuration guides and best practices
	Integration	FIM software may support integration with other reporting and alerting capabilities. Evaluate: • Supported plugins and integrations with business and security platforms in use by the organization (such as AWS, ticketing, SIEM, incident response, threat intelligence) and the capabilities of these plugins and integrations • API support (including API-first, REST API available, programmatic API available) • Whether the platform allows the customer to build custom plugins or integrations • Level of effort required and technology (languages and frameworks) supported when building custom plugins or integrations
	Reporting, metrics and alerting	If a monitored file is changed, human intervention is typically required to determine the cause and whether it was an approved change. Adequate reporting and alerts are needed to facilitate this process. Evaluate: • Out-of-the-box reports and dashboards against current program requirements • Ability and level of effort required to create custom measures and metrics • Alerting mechanisms and ability to create or modify alerts • Supported reporting and alerting formats and delivery mechanisms • Integration with AWS reporting and alerting tools (such as AWS Security Hub, Amazon CloudWatch Events, Amazon SNS) • Resources and processes to support monitoring of reports and response to alerts • Support for data aggregation across regions • Supported data export formats

FIM is one of the easier technologies to implement for most organizations. Depending on the configuration, however, the number of files being monitored and the amount of change in the organization, the number of resources required to follow up on alerts can be excessive. Continuous tuning and integration with change management can help reduce to a manageable level the number of alerts requiring human interaction.

Application Whitelisting

Application whitelisting protects endpoints by either ensuring that only known software is allowed to execute or notifying administrators when unapproved software is executed on endpoints. This protection may be accomplished by validating hashes or signatures associated with the software or by validating software-signing certificates against the policies defined by the organization and assigned to each endpoint. Application whitelisting makes the exploitation and installation phases of the attack kill chain much more difficult. Consider application whitelisting if your environment has a high degree of homogeneity or if your organization's deployment processes are mature and would support automating the development and maintenance of the whitelist policies.

Caution: Application whitelisting technologies may not prevent attacks against known vulnerabilities in whitelisted applications, so be sure to follow good vulnerability management practices.

	Consideration	Details
	Performance and efficiency	Application whitelisting solutions generally do not affect performance as much as other endpoint security solutions—as long as they are configured correctly. If they are misconfigured and block legitimate applications or services, the performance impact is significant. Changes to rules and to cloud resources should be thoroughly tested before deployment. Evaluate: • The architecture of the tools under consideration • Performance (CPU, memory, storage and bandwidth utilization) when used with production workloads • Amount of data sent and received from management console(s) • Support for data compression
	Deployment	Application whitelisting solutions typically require agents and a management console to update and distribute configurations and receive alerts from agents. Evaluate: • The architecture of the tools under consideration • The installation and configuration procedures for agents and management infrastructure • The availability of managed or SaaS components or preconfigured appliances from AWS Marketplace • Effectiveness and responsiveness of support • Any vendor requirements for the use of professional services for installation or configuration • Integration with other AWS technologies (such as AWS Systems Manager, AWS Config, Amazon CloudWatch) for deployment or validation of agent deployment

Application Whitelisting (continued)

	Consideration	Details
	Configuration and maintenance	Configuration and maintenance of whitelisting policies is critical to the successful use of application whitelisting. In enterprise environments, standardization and automation can help reduce this burden. Automated testing can validate changes to the whitelist or cloud resources before release into production environments. Evaluate: • The architecture of the tools under consideration • The upgrade procedures • The procedures for updating whitelists • Reporting, metrics or alerting available for any out-of-date agents or policies • Communication protocols and paths to understand required firewall and ACL changes along with any VPC peering or cross-account access • Any vendor requirements for the use of professional services for upgrades or updates • Level of effort to create and maintain whitelists and any assistance provided by technology • Integrations with other AWS technologies (such as AWS Config or AWS Lambda) or configuration management tools (Puppet, Chef, Ansible, SaltStack, CFEngine) to perform updates and upgrades or to apply configurations • Secure configuration guides and best practices
	Reporting, metrics and alerting	In order to respond quickly to outages caused by whitelists and aid in the identification of attempted exploits and unauthorized installations, evaluate the reporting and alerting features available. Evaluate: • Support for centralized logging technologies and communication protocols including integration with any existing or proposed SIEM technology • Out-of-the-box reports and dashboards against current program requirements • Ability and level of effort required to create custom measures and metrics • Alerting mechanisms and ability to create or modify alerts • Supported reporting and alerting formats and delivery mechanisms • Integration with AWS reporting and alerting tools (such as AWS Security Hub, Amazon CloudWatch Events, Amazon SNS) • Support for data aggregation across regions • Supported data export formats

Application whitelisting is a mature, layered security control that can be leveraged to reduce the impact of vulnerabilities in cloud environments and make exploitation of cloud resources more difficult. Because standardization is more common in the cloud, application whitelisting may be more achievable and easier to maintain. Heavy use of automation and DevOps principles can also help ease the burden of ongoing configuration and maintenance.

Making the Choice

To summarize, the key considerations for implementing endpoint security in AWS are:

- Cloud context
- Efficiency
- Ease of use
- Reporting
- Ease of integration
- Effectiveness

Have a Plan

By defining and understanding their cloud architecture, risk profile, business requirements and available resources along with understanding any gaps, organizations will be in a good position to determine which considerations outlined above are most important to them. Based on those considerations, organizations should develop a proof-of-concept test plan and evaluation matrix. The test plan and matrix should include a ranking of importance for each consideration, and where possible, acceptance thresholds. When the test plan is complete, the organization should identify two or more representative cloud environments in which to conduct the test. They should identify any additional technology they may need to aid in the evaluation of certain considerations. For example, evaluating the performance and efficiency of agents will most likely require additional setup and configuration, and, depending on the platform, performance monitoring tools may be required.

Consider Partners

As organizations build out their cloud and cloud security strategy and plan, they may want to consider working with partners to accelerate their efforts or fill any gaps in knowledge or resources that are identified. All consulting partners may extend AWS Marketplace third-party solutions directly to customers through Consulting Partner Private Offers (CPPO). Not every organization will be able to find resources with deep cloud experience and even experienced cloud technologists may only have experience in specific industries or with specific cloud vendors.

Test and Evaluate

With the plan and any additional requirements in place, the technology should be installed in the test environment, configured and monitored to gather enough data to evaluate each consideration. Every step of the process should be measured.

Organizations, if possible, should avoid allowing vendors to install and configure the technology for the proof of concept unless they will be installing and managing the solution after purchase as well. At a minimum, technical resources should be available to observe these processes.

After the proof-of-concept test, organizations should evaluate the results against the test plan and acceptance thresholds. Use the collected and documented results to compare functionality, cost and other factors to determine the best solution(s) to employ.

Conclusion

Endpoint security for IaaS cloud workloads is an important part of an organization's cloud security strategy. Not only does it provide additional protections for these workloads, but it also provides additional visibility into cloud resources and the actual threats that exist in an organization's cloud environments. While many organizations are still concerned about the performance impacts and associated costs, cloud endpoint security vendors have matured, and cloud-optimized solutions are more accessible.

Fortunately, many of these solutions are offered on-demand, which makes evaluating these products and services much easier than it was in the past. To get started, you may want to review what products are available in AWS Marketplace or through a SaaS model to jump-start your evaluation process.

About the Author

David Hazar is a SANS analyst, instructor and co-author of SANS MGT516: Managing Security Vulnerabilities: Enterprise and Cloud. He also is an instructor for SANS SEC540: Cloud Security and DevOps Automation. With close to 20 years of broad, deep technical experience gained from a variety of hands-on roles serving the financial, healthcare and technology industries, his current areas of focus include vulnerability management, application security, cloud security and secure DevOps. He holds the CISSP, GWAPT, GWEB, GMOB, GCIA, GCIH, GCUX, GCWN, GSSP-.NET and GSTRT certifications.