



The most trusted source for
cybersecurity training, certifications,
degrees, and research

Cybersecurity in the Age of the Cloud

February 2020

Compiled from works completed by

Kyle Dickinson

Shaun McCullough

John Pescatore

Dave Shackleford

David Szili

with a forward by

Frank Kim

Table of Contents

3 **Forward by Frank Kim**

5 **Additional Cloud Resources**

6 **How to Optimize Security Operations in the Cloud Through the Lens of the NIST Framework**

by John Pescatore

Security teams today face the mandate of moving production workloads from on-premises to the cloud. By using the NIST Cybersecurity Framework (CSF), teams can effectively and efficiently build in security as part of the migration of operations activities to IaaS services and hybrid cloud implementations. This paper shares proven best practices for evaluating and implementing security architectures, processes and controls while developing an approach to migration that is repeatable.

21 **How to Build a Data Security Strategy in AWS**

by Dave Shackleford

When organizations move sensitive data to the cloud, they absolutely must choose a provider that can ensure compliance with privacy regulations on a global stage. Data security strategies in the cloud must include encryption and key management, data loss prevention, and the capability to classify and track data. By using the AWS Cloud, organizations can protect sensitive data at rest, in transit, and in use.

37 **How to Protect a Modern Web Application in AWS**

by Shaun McCullough

In moving assets to the cloud, organizations need to prioritize their security plans based on the risks to which they are exposed. With threat modeling, organizations can identify and prioritize the risks to infrastructure, applications, and the services they provide, as well as evaluate how to manage those risks over time. This paper includes use cases for threat modeling web apps and the DevSecOps platform, using a process that is both repeatable and improvable.

51 **How to Build a Threat Detection Strategy in Amazon Web Services (AWS)**

by David Szili

Threat detection and continuous security monitoring in the cloud must integrate traditional on-premises system monitoring with the cloud network infrastructure and cloud management plane. A successful, cloud-based threat detection strategy will collect data from systems, networks, and the cloud environment in a central platform for analysis and alerting. This paper describes how to build a threat detection strategy that automates common tasks like data collection and analysis.

68 **How to Build a Threat Hunting Capability in AWS**

by Shaun McCullough

Threat hunting is more of an art than a science, in that its approach and implementation can differ substantially among enterprises and still be successful. In cloud environments, where the threat landscape is always changing, security teams must know what data to collect and how to analyze it in order to tease out suspicious anomalies. In addition to these topics, this whitepaper walks you through the threat hunting process, describing tools and techniques you can use to find and neutralize threats.

84 **How to Perform a Security Investigation in AWS**

by Kyle Dickinson

Because the technologies that enable investigations in the cloud differ from those on premises, as do the levels of responsibility, organizations need to put in place a cloud-specific incident response plan. By planning out how they will perform investigations using solutions such as AWS, organizations can validate that any obligations they may have as a security organization can be met as effectively in cloud environments as they did in-house.

92 **SANS 2019 Cloud Security Survey**

by Dave Shackleford

This whitepaper delves into the results of the SANS 2019 Cloud Security Survey, conducted in cooperation with the Cloud Security Alliance, which covers organizations' use of the public cloud and provides actionable advice organizations can use to improve their cloud security. It answers questions including, "Are security infrastructures maturing to support the business and improve risk management in the cloud model?" and "How are organizations using the public cloud to meet their business needs?"

Forward

Just as the development of the Internet defined new challenges and principles for information security 30 years ago, cloud computing has become a defining factor in the current state and future of information security.

Every organization, large and small, is facing a new reality of mandatory, multi-cloud integrations and management. Even organizations that previously expressed security reservations about migrating to the cloud seem to be gaining confidence in cloud providers and are dipping their toes into the cloud services pool. Data from the SANS 2019 Cloud Security Survey bears that out: Just 44% of respondents indicated concern for data breaches by cloud provider personnel, down 8% from those expressing the same concern in SANS' 2017 survey.

By Frank Kim
SANS Senior
Instructor and
Curriculum Lead

Whether or not you interpret this data point as a sign of growing trust in cloud providers, the truth is that the business reasons for moving to the cloud are simply too overwhelming to ignore. Besides the often-touted benefits of cost savings and elasticity, one of the big concepts of the cloud is that it makes automation substantially easier compared to the pre-cloud environment where people had to set up their own duplicative infrastructures to spin things up. It's when DevOps and cloud are leveraged in tandem that organizations truly take full advantage of the cloud and its promise to change for the better the way that people work.

The cloud represents big change for almost all organizations, and security must be part of that evolution to succeed. If you haven't already begun building out your cloud security knowledge base, SANS is here to help.

So, what types of things do you need to know and how can SANS help you advance your cloud capabilities and policies? Today's security professionals need to do three things in parallel:

- You need to understand how the major cloud providers work and the plethora of services that they offer.
- You must understand the technical details of each platform to ensure that you have secured your specific implementation appropriately.
- And finally, security teams must transform the way that we work to leverage cloud services and automation to improve the effectiveness of security itself.

SANS

The hand-selected resources in this eBook provide a well-rounded look at cybersecurity considerations and practices in the age of the cloud. Each report touches on different parts of the five functions of the NIST Cybersecurity Framework – identify, protect, detect, respond, and recover. Together, the collection models the whole lifecycle of security.

The eBook is rounded out by the recent SANS 2019 Cloud Security Survey to provide a snapshot of today’s cloud security environment and associated concerns.

In terms of industry momentum, we’ve reached the point where every cybersecurity professional has to be knowledgeable about the cloud to varying degrees. Use these resources to carve out a path for your career and your organization in this new cloud security future.



Frank Kim
SANS Senior Instructor
[@fykim](#)

Frank is the founder of ThinkSec, a security consulting and CISO advisory firm. Previously, as CISO at the SANS Institute, Frank led the information risk function for the most trusted source of computer security training and certification in the world. Frank continues to lead the management and software security curricula at SANS, helping to develop the next generation of security leaders. Frank was also executive director of cybersecurity at Kaiser Permanente, where he built an innovative security program to meet the unique needs of the nation’s largest not-for-profit health plan and integrated health care provider with annual revenue of \$60 billion, 10 million members, and 175,000 employees. Frank holds degrees from the University of California at Berkeley and is the author and instructor of popular courses on strategic planning, leadership, application security, and DevOps.



Get More on Cloud Security From SANS

Not only can you get cloud-focused cybersecurity training from SANS, but you can validate your skills with a GIAC Certification, deepen your cloud connections with Summits, and expand your cloud knowledge base with tons of [free resources](#) – webcasts, blogs, tools, research, and other resources.



Cloud Security Training and Certifications

This rapidly developing focus area incorporates hands-on labs and exercises to help cybersecurity professionals apply their skills within real-world cloud environments.

[SEC540: Cloud Security and DevOps Automation](#)

[SEC545: Cloud Security Architecture and Operations](#)

[MGT516: Managing Security Vulnerabilities: Enterprise and Cloud](#)

[GCSA Cloud Security Automation Certification](#)



Cloud Security Blogs

Enjoy a selection of recent blogs on cloud security topics:

[Breaches in the Cloud and Why Blame Matters](#)

[Security Visibility in AWS: Authority, Access and Capability](#)

[Workforce Automation and the Cloud: A Dichotomy](#)

[How to Build a Successful Cloud Security Program](#)



Cloud Security Webcasts

SANS hosts hundreds of webcasts per year, many touching on topics related to cloud security, like:

[Cloud Security Vulnerabilities, Management, and Communication](#)

[Getting Your Head in the Cloud: Implementing Security Monitoring for Public Clouds](#)

[Cloud Storage Forensics: Endpoint Evidence](#)

[SEC588: Cloud Penetration Testing. What is it? What's different, and why?](#)

Cloud Security Summit & Training

Frisco, TX | May 27 – June 2

SANS

sans.org/u/ZMp

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

Written by **John Pescatore**

February 2019

Sponsored by:

AWS Marketplace

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.

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.

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.

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 to 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.

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

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.
- **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

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.

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.

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

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.

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 CSF⁶

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/

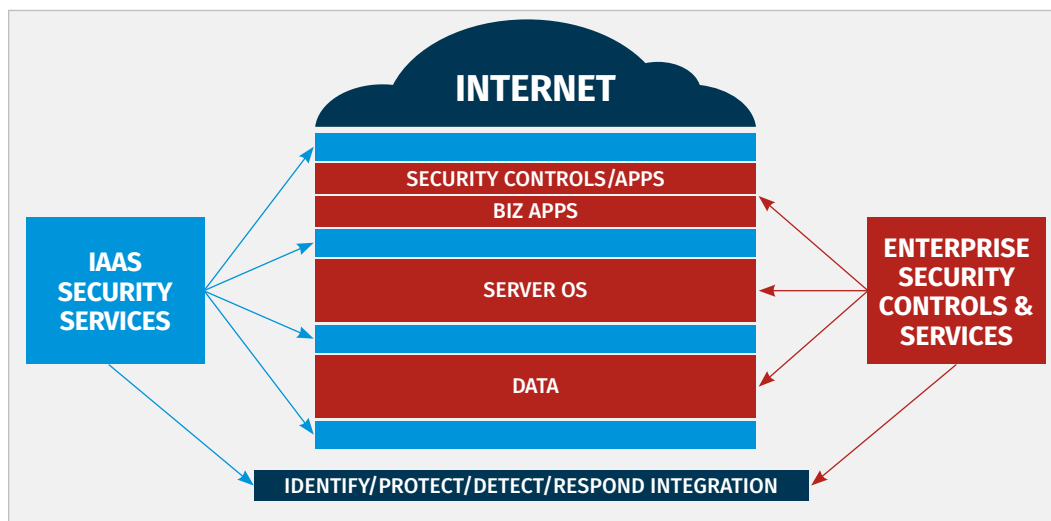


Figure 2. Integrated Services at Each Boundary Layer

analysis would be considered noncompliant. While compliant does not always mean secure, noncompliant almost always represents unacceptable business risk.

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 risk exposure. 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.

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

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.

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
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 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 on the next page.

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.

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.

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
Reactive	Detect	Ops skills training	Firewall policy management
		Intrusion detection systems	Data analytics
		Network monitoring	Data loss prevention
	Respond	SIEM	
		Incident response services	Endpoint detect/respond
	Recover	Trouble ticket systems	Forensic analysis
		System/endpoint backup	High-avail/mirroring services

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?

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

- **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?

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.

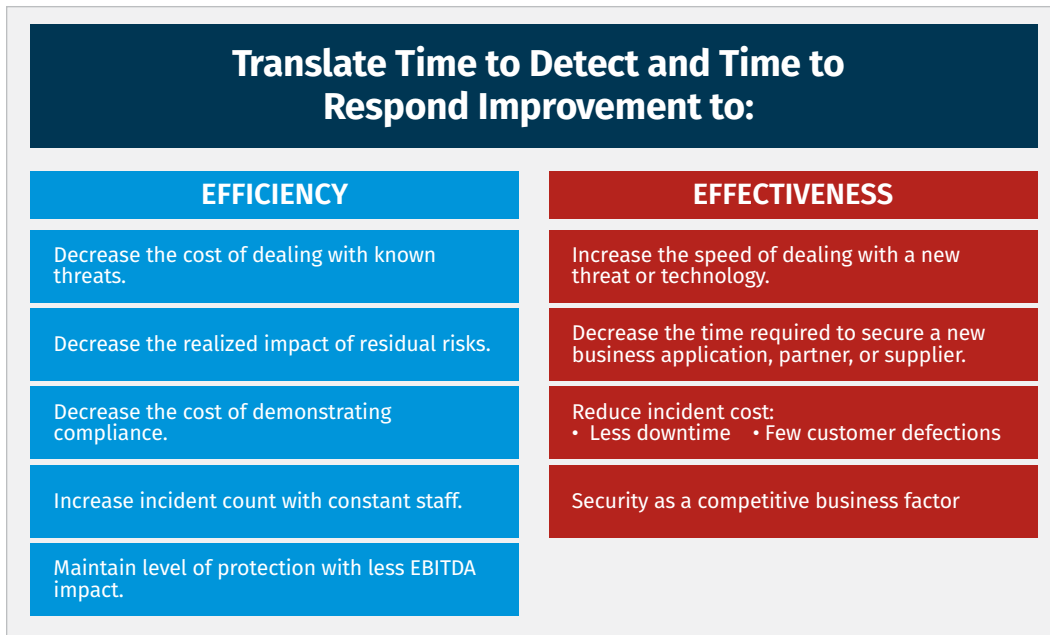


Figure 3. Connecting Metrics to Business Services

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.

Sponsor

SANS would like to thank this paper’s sponsor:



References

- 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
- 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
- ⁴ “National Institute of Standards and Technology,” www.nist.gov/news-events/news/2018/04/nist-releases-version-11-its-popular-cybersecurity-framework
- ⁵ Cybersecurity Enhancement Act of 2014, www.congress.gov/bill/113th-congress/senate-bill/1353/text
- 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>
- ⁷ www.gartner.com/document/500890?ref=solrResearch&refval=214539204&qid=d3f5b689a39463b6c77406155a9672a1 [Registration required for access.]
- ⁸ Risk Management, NIST, [https://csrc.nist.gov/projects/risk-management/risk-management-framework-\(rmf\)-overview](https://csrc.nist.gov/projects/risk-management/risk-management-framework-(rmf)-overview)
- ⁹ CIS Benchmarks, Center for Internet Security, www.cisecurity.org/cis-benchmarks

How to Build a Data Security Strategy in AWS

Written by **Dave Shackelford**

May 2019

Sponsored by:

AWS Marketplace

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.

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.

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 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.

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.

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.

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.

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.

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.

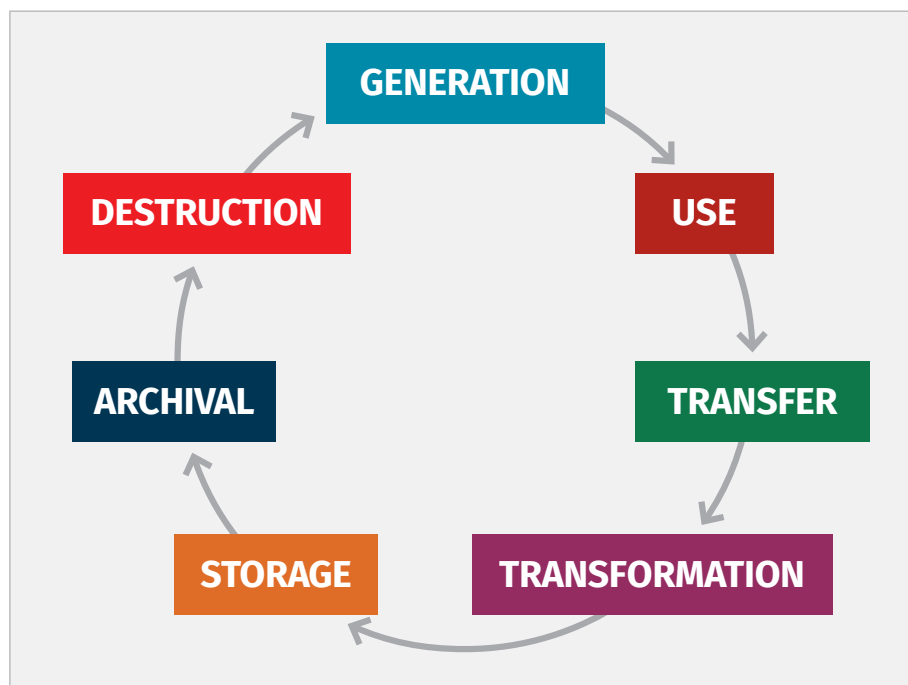


Figure 1. Data Life Cycle Model

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.

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.

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.

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.

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.

Sponsor

SANS would like to thank this paper's sponsor:



References

- 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
- 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.

How to Protect a Modern Web Application in AWS

Written by **Shaun McCullough**

May 2019

Sponsored by:

AWS Marketplace

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.

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.

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.

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

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.

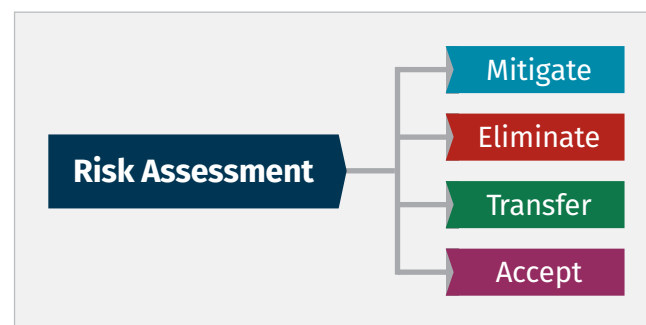


Figure 1. Risk Management Strategies

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

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.

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. Security teams that work separately from those who create the systems are fighting an uphill battle that will impair effectiveness while raising costs. 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.

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

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.

DevOps with Security

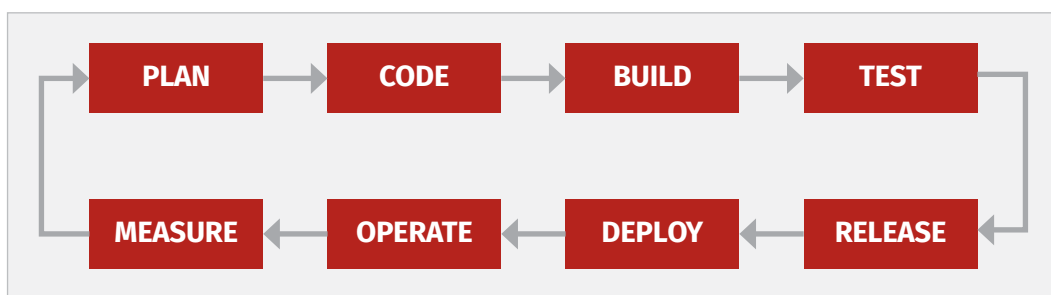
DevOps is a process that enables close coordination between development and operation teams.³ That integration enables organizations to develop and quickly deploy new services with zero downtime and improved reliability. The process is especially beneficial for organizations that deploy new versions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

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

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.

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 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 2 on the next page describes the ratings.

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 2. DREAD Risk Ratings⁷

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	

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.⁸ 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.

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

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 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.

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 3.

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.

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 on the next page.

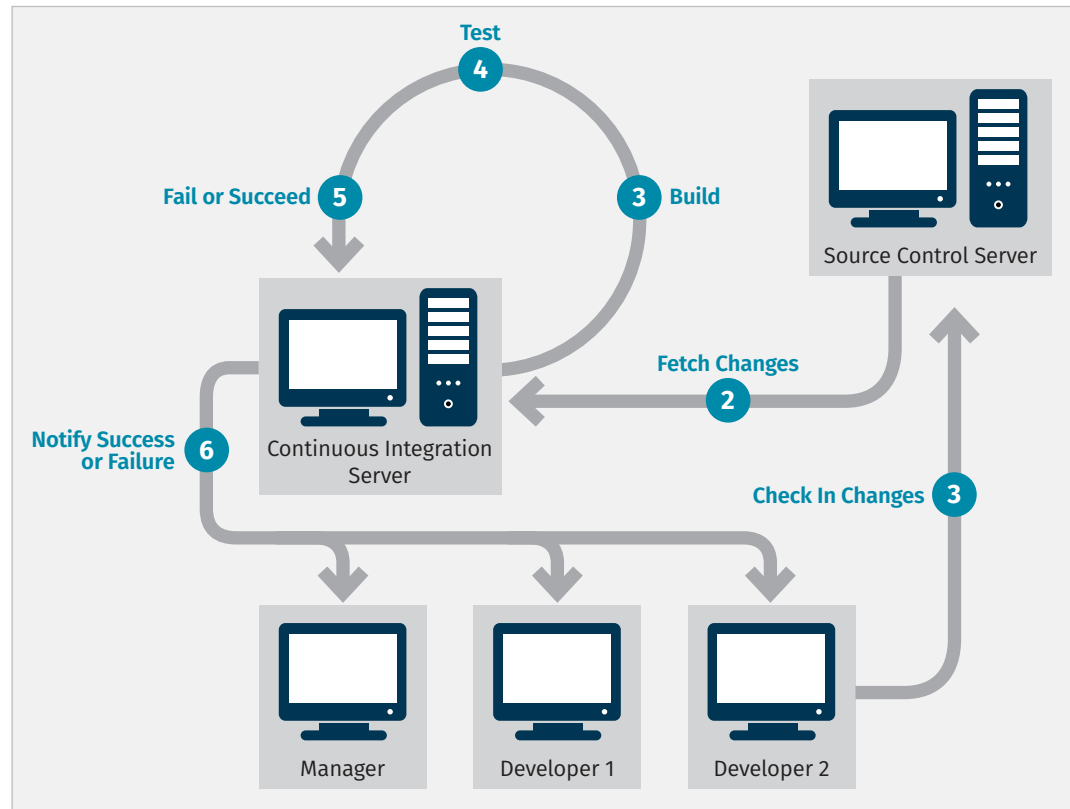


Figure 3. Continuous Integration Process

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	

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.

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.

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	

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.

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 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

Shaun McCullough is a community instructor for the [SEC545: Cloud Security Architecture and Operations](#) class and gives back to his profession by mentoring and supporting the next generation of cyber professionals. With 25 years of experience as a software engineer, he has been focusing on information security for the past 15 years. Shaun is a consultant with H&A Security Solutions, focusing on secure cloud operations, building DevSecOps pipelines and automating security controls in the cloud. He also served as technical director of red and blue team operations, researched advanced host analytics, and ran threat intelligence on open source platforms in his work with the U.S. Department of Defense.

Sponsor

SANS would like to thank this paper's sponsor:



References

- ¹ Draft NIST SP 800-154, Guide to Data-Centric System Threat Modeling, <https://csrc.nist.gov/publications/detail/sp/800-154/draft>
- ² IANS Pragmatic Threat Modeling Toolkit, https://portal.iansresearch.com/media/739278/ians_pragmatic_threat_modeling_toolkit.xlsm
- ³ NIST SP800-190, <https://nvlpubs.nist.gov/nistpubs/specialpublications/nist.sp.800-190.pdf>
- ⁴ *Accelerate: Building and Scaling High Performing Technology Organizations*, by Nicole Forsgren, Jez Humble and Gene Kim (IT Revolution, 2018)
- ⁵ OWASP Top Ten Project, www.owasp.org/index.php/Category:OWASP_Top_Ten_Project
- ⁶ The 10 Most Critical Web Application Security Risks, [www.owasp.org/images/7/72/OWASP_Top_10-2017_\(en\).pdf](http://www.owasp.org/images/7/72/OWASP_Top_10-2017_(en).pdf)
- ⁷ Adapted from DREAD Rating, <https://wiki.openstack.org/wiki/Security/OSSA-Metrics#DREAD>
- ⁸ Web Application Firewall, www.owasp.org/index.php/Web_Application_Firewall
- ⁹ This 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.
- ¹⁰ SQL Injection: Modes of Attack, Defence, and Why It Matters, www.sans.org/reading-room/whitepapers/securecode/sql-injection-modes-attack-defence-matters-23
- ¹¹ SQL Injection Bypassing WAF, www.owasp.org/index.php/SQL_Injection_Bypassing_WAF
- ¹² Getting Started—About Version Control, <https://git-scm.com/book/en/v2/Getting-Started-About-Version-Control>
- 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>
- ¹⁵ NPM security advisories, www.npmjs.com/advisories

How to Build a Threat Detection Strategy in Amazon Web Services (AWS)

Written by **David Szili**

August 2019

Sponsored by:

AWS Marketplace

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



Figure 1. Steps to Build a Security Monitoring Strategy

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.

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 on the next page shows that a root user

```
("userIdentity": { "type": "Root"})  
successfully signed into the AWS Management Console  
("eventName": "ConsoleLogin") using multifactor authentication  
("MFAUsed": "Yes");
```

```

{
  "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

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.

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.



Figure 3. Unified Interface for Management of Events and Alerts

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.

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.

By reducing the amount of manual labor in the team, the team has more time to focus on other areas of information security.

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.

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.
 - Billing alarms
 - 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.

Sponsor

SANS would like to thank this paper's sponsor:



References

¹ Because this paper is an exploration of threat detection 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.

How to Build a Threat Hunting Capability in AWS

Written by **Shaun McCullough**

November 2019

Sponsored by:

AWS Marketplace

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 threats and undetected breaches.

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 victim? 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.

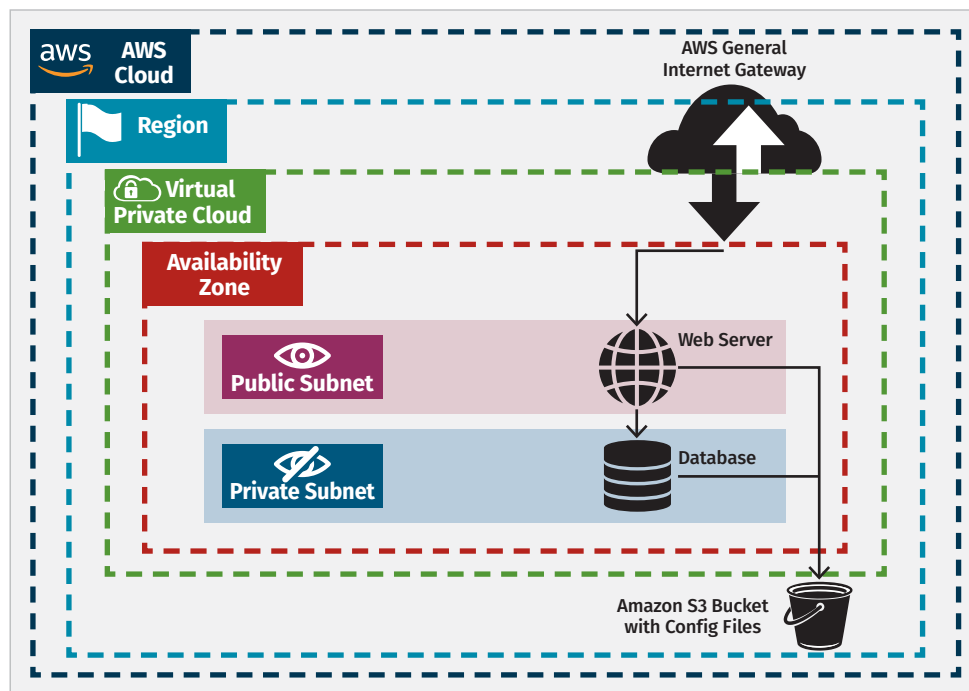


Figure 1. Web Application Use Case

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.

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 server administration 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.

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.

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 Hunting²

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 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.

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.

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 and security companies (see Figure 3 on the next page). Although not cloud-specific, the ATT&CK Framework provides a detailed explanation of the hows and whys of specific attacker techniques.



Figure 2. Threat Hunting Loop

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	Valid Accounts		Forced Authentication	Network Share Discovery		AppleScript	Man in the Browser	Exfiltration Over Physical Medium	Multi-Hop Proxy
Plist Modification	Valid Accounts		Hooking	System Time Discovery		Third-party Software	Browser Extensions	Medium	Domain Fronting
DLL Search Order Hijacking	Valid Accounts		Password Filter DLL	Peripheral Device Discovery		Windows Remote Management	Video Capture	Exfiltration Over Command and Control Channel	Data Encoding
AppCert DLLs	Process Doppelgänger		LLMNR/NBT-NS Poisoning	Account Discovery	SSH Hijacking	LSASS Driver	Audio Capture	Data Encrypted	Remote File Copy
Hooking	Mshta		Securityd Memory	File and Directory Discovery	Distributed Component Object Model	Dynamic Data Exchange	Automated Collection	Scheduled Transfer	Multi-Stage Channels
Startup Items	Hidden Files and Directories		Private Keys	System Information	Pass the Ticket	Local Job Scheduling	Clipboard Data	Data Encrypted	Web Service
Launch Daemon	Launchctl		Keychain	Discovery	Replication Through Removable Media	Source	Screen Capture	Exfiltration Over Other Network Medium	Standard Non-Application Layer Protocol
Dylib Hijacking	Space after Filename		Input Prompt	Security Software Discovery	Windows Admin Shares	Launchctl	Trap	Exfiltration Over Alternative Protocol	Communication Through Removable Media
Application Shimming	LC_MAIN Hijacking		Bash History	System Network Connections Discovery	Remote Desktop Protocol	Space after Filename	Data from Network Shared Drive	Data Transfer Size Limits	Standard Application Layer Protocol
Appinit DLLs	HISTCONTROL		Two-Factor Authentication Interception	System Owner/User Discovery	Pass the Hash	Execution through Module Load	Data from Local System	Data Compressed	Commonly Used Port
Service Registry Permissions Weakness	Hidden Users		Account Manipulation	System Network Configuration	Exploitation of Vulnerability	Regsvcs/Regasm	Data from Removable Media		Standard Cryptographic Protocol
Scheduled Task	Gatekeeper Bypass		Replication Through Removable Media	Discovery	Shared Webroot	InstallUtil			Custom Cryptographic Protocol
New Service	Hidden Window		Input Capture	Application Window Discovery	Logon Scripts	Regsvr32			Data Obfuscation
File System Permissions Weakness	Deobfuscate/Decode Files or Information		Network Sniffing	Discovery	Remote Services	Regsvr32			Custom Command and Control Protocol
Path Interception	Trusted Developer Utilities		Credential Dumping	Network Service Scanning	Application Deployment Software	PowerShell			Connection Proxy
Accessability Features	Regsvcs/Regasm		Brute Force	Query Registry	Remote File Copy	Rundll32			Uncommonly Used Port
Port Monitors	Exploitation of Vulnerability		Credentials in Files	Remote System Discovery	Taint Shared Content	Scripting			Multi-band Communication
Screensaver	Extra Window Memory Injection			Permission Groups Discovery		Graphical User Interface			Failback Channels
LSASS Driver	Access Token Manipulation			Process Discovery		Command-Line Interface			
Browser Extensions	Bypass User Account Control			System Service Discovery		Scheduled Task			
Local Job Scheduling	Process Injection					Windows Management Instrumentation			
Re-opened Applications	SID-History Injection	Component Object Model Hijacking				Trusted Developer Utilities			
Rc-common	Sudo	InstallUtil				Service Execution			
Login Item	Setuid and Setgid	Regsvr32							
LC_LOAD_DYLIB Addition	Launch Agent	Code Signing							
Hidden Files and Directories		Modify Registry							
.bash_profile and .bashrc		Component Firmware							
Trap		Redundant Access							
Launchctl									

Figure 3. MITRE ATT&CK Framework⁵

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 on the next page.) 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.

ENTERPRISE ▾
Home > Techniques > Enterprise > Exploit Public-Facing Application

TECHNIQUES

All

- Initial Access -
- Drive-by Compromise
- Exploit Public-Facing Application
- External Remote Services
- Hardware Additions
- Replication Through Removable Media
- Spearpishing Attachment
- Spearpishing Link
- Spearpishing via Service
- Supply Chain Compromise
- Trusted Relationship
- Valid Accounts
- Execution +
- Persistence +
- Privilege Escalation +
- Defense Evasion +
- Credential Access +
- Discovery +
- Lateral Movement +

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]}

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

ID: T1190

Tactic: Initial Access

Platform: Linux, Windows, macOS

Data Sources: Packet capture, Web logs, Web application firewall logs, Application logs

Version: 1.1

Figure 4. The Exploit Public-Facing Application Technique⁶

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.

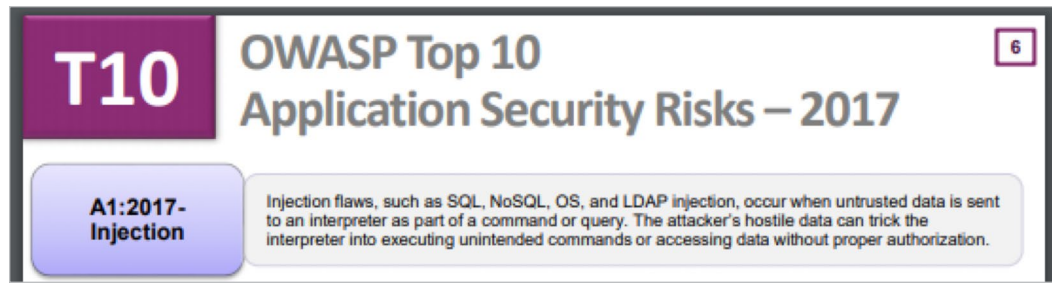


Figure 5. Number One Threat in the OWASP Top 10⁷

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.

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.

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

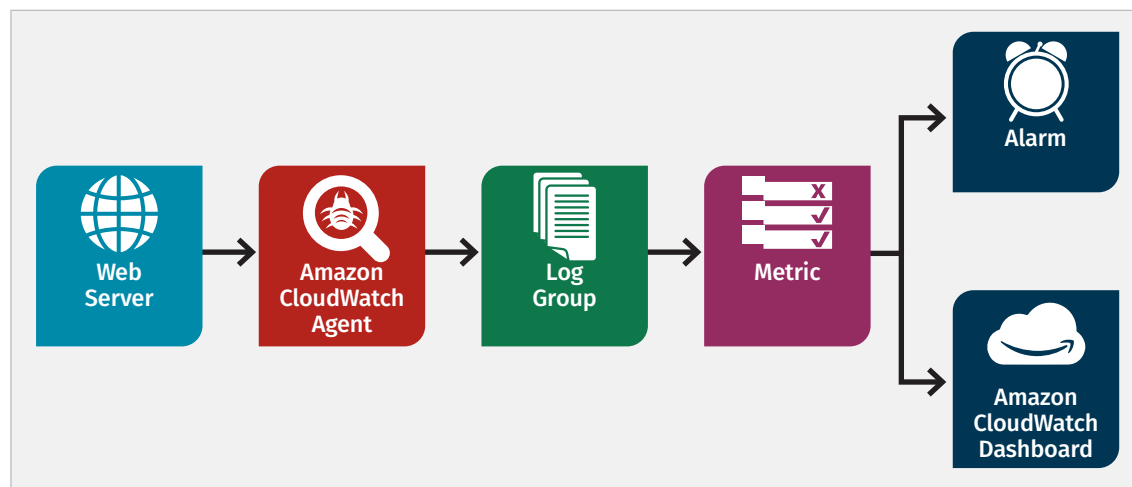


Figure 6. Overview of Amazon CloudWatch Log Collection

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	Time range:		
	i-0b1515ec2d4b0b9df	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[0].{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.

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.

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

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.

```
cybergooft> aws cloudtrail lookup-events --lookup-attribute AttributeKey=EventI
d,AttributeValue=396ef72d-8b25-4adc-a84a-7f0e4a09be3f --query "Events[0].CloudT
railEvent" --region us-east-1 --output text | jq -r '.'
{
  "eventVersion": "1.05",
  "userIdentity": {
    "type": "AssumedRole",
```

Figure 10. JSON Output of AWS CloudTrail lookup-events

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

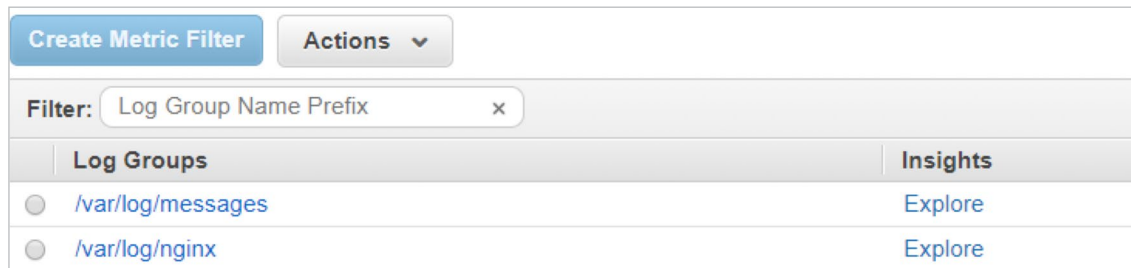
For the web application use case, the VM itself has a wealth of information that could be of interest. Mainstream web servers generate standard logs that are stored on the VM. They also can be customized to generate more or fewer logs, or with changes to the format or location, and potentially compressed 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 attack¹⁰ 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.

The Nginx connection logs are now stored in the `/var/log/nginx` log group, accessible from Amazon CloudWatch Logs. See Figure 12.

```
[/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



Log Groups		Insights
<input type="radio"/>	/var/log/messages	Explore
<input type="radio"/>	/var/log/nginx	Explore

Figure 12. Nginx Connection Logs

passwd	
Time (UTC +00:00)	Message
2019-09-15	
<i>No older events found for the selected filter and date range. Adjust the date range or clear filter.</i>	
▶ 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; rv:68.0) Gecko/20100101 Firefox/68.0"
▶ 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; rv:68.0) Gecko/20100101 Firefox/68.0"
▶ 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; rv:68.0) Gecko/20100101 Firefox/68.0"
▶ 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; rv:68.0) Gecko/20100101 Firefox/68.0"

Figure 13. Quick Search for passwd

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 `\`.

```
fields @timestamp, @message
| sort @timestamp desc
| limit 20
| filter @message like /passwd|etc|..\/|
```

Figure 14. Query Amazon CloudWatch Logs Insights

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.

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.

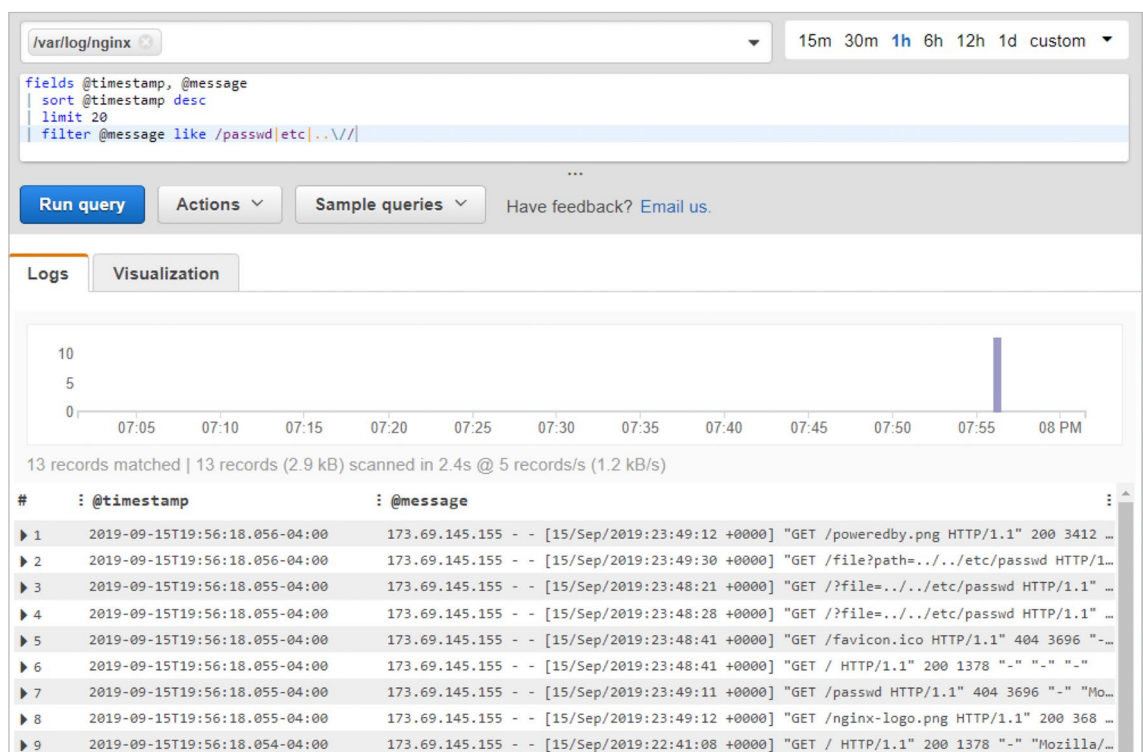


Figure 15. Query Results

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.

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.

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).

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.

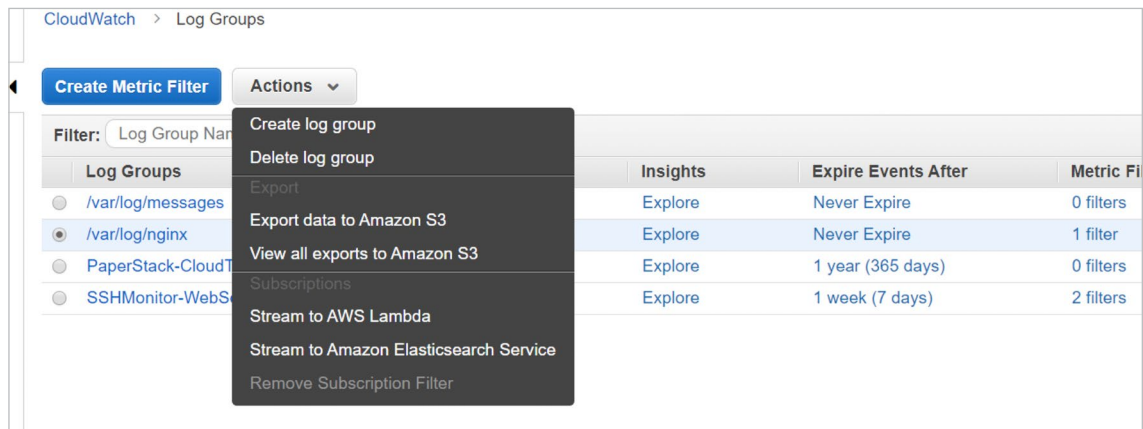


Figure 16. Exporting Amazon CloudWatch Logs

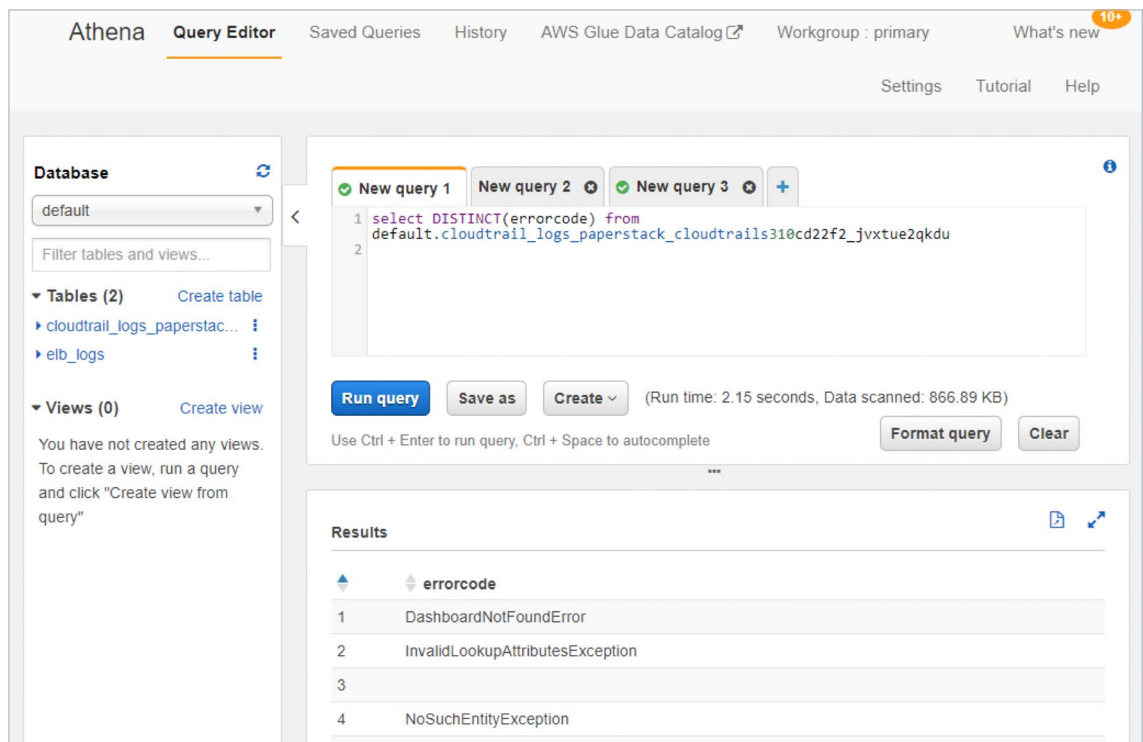


Figure 17. Amazon Athena Dashboard

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.

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

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 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 Palo Alto's Demisto and 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.

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. The SOAR then performs a memory dump of the EC2, takes a snapshot and stops the EC2. All the data is gathered up and prepared in an Amazon S3 bucket for the security team when it is ready to investigate.

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

Shaun McCullough is a SANS instructor for the [SEC545: Cloud Security Architecture and Operations](#) class and gives back to his profession by mentoring and supporting the next generation of cyber professionals. With 25 years of experience as a software engineer, he has been focusing on information security for the past 15 years. Shaun is a consultant with H&A Security Solutions, focusing on secure cloud operations, building DevSecOps pipelines and automating security controls in the cloud. He also served as technical director of red and blue team operations, researched advanced host analytics, and ran threat intelligence on open source platforms in his work with the U.S. Department of Defense.

Sponsor

SANS would like to thank this paper's sponsor:



References

- 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.
- ² www.cisecurity.org/controls/cis-controls-list/
- ³ www.threathunting.net/sqrrl-archive
- 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.]
- ⁵ <https://attack.mitre.org/>
- 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
- ⁸ Cloud Security Alliance, Top Threats to Cloud Computing: Egregious Eleven, <https://cloudsecurityalliance.org/artifacts/top-threats-to-cloud-computing-egregious-eleven/>
- ⁹ Cloud Security Alliance, Top Threats to Cloud Computing: Deep Dive, <https://cloudsecurityalliance.org/artifacts/top-threats-to-cloud-computing-deep-dive/>
- ¹⁰ www.owasp.org/index.php/Path_Traversal

How to Perform a Security Investigation in AWS

Written by **Kyle Dickinson**

October 2019

Sponsored by:

AWS Marketplace

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.

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

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).

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 steps for responding using the SANS six-step incident response methodology.³



Figure 1. SANS Incident Response Steps

This methodology can easily be adapted to cloud-based environments. 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.

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

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

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.

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.

```

{
  "eventVersion": "1.05",
  "userIdentity": {
    "type": "IAMUser",
    "principalId": "AIDAZDEVHULLOJ65ACNU",
    "arn": "arn:aws:iam:90123456789:user/Marc_the_Intern",
    "accountId": "90123456789",
    "userName": "Marc_the_Intern"
  },
  "eventTime": "2019-09-04T23:00:13Z",
  "eventSource": "signin.amazonaws.com",
  "eventName": "ConsoleLogin",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "11.22.33.44",
  "userAgent": "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_15_0; rv:68.0) Gecko/20100801 Firefox/68.0",
  "requestParameters": null,
  "responseElements": {
    "ConsoleLogin": "Success"
  },
  "additionalEventData": {
    "LoginTo": "https://console.aws.amazon.com/console/home",
    "MobileVersion": "No",
    "MFAUsed": "No"
  },
  "eventID": "734f86de-ff17-47ef-8e60-5e6186fe041d",
  "eventType": "AwsConsoleSignIn",
  "recipientAccountId": "90123456789"
}

```

The **userIdentity** used for the event:

- type**: Shows if a role or user was used
- principalId**: Unique identifier for this specific user (Think SID)
- arn**: Amazon Resource Name
- accountId**: Which account ID was logged into
- userName**: User that authenticated

Additional details:

- eventTime**: Zulu time for when the event occurred
- eventSource**: How the API was called
- eventName**: One of many API calls that can be used within AWS
- awsRegion**: Which region the console was set to log into (can vary depending on how the login was initiated; good source to determine if activity is occurring outside of normal regions)
- sourceIPAddress**: The IP address that the request was sent from
- userAgent**: Fingerprint of what was used (browser or CLI version)
- requestParameters**: What was included in the request
- responseElements**: If the API delivers a response, this section contains additional details

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.




Figure 3. Structure of a VPC Flow Log

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.

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.

Use Case: An Investigation

Consider a case where the internal audit organization has approached the security organization. The audit organization requires an investigation of the user, Marc the Intern. It also requests that the security team acquire a forensic image, summarize that image and include a summary of the communications the instances had if Marc created any Amazon EC2 instances.

With running the Amazon EC2 instance, the security team wants to understand what this instance is doing so it can perform further analysis. After acquiring a snapshot, the team converts the snapshot to a volume so that it may attach the new volume that contains evidence to its analysis instance.

The team finds that Marc had access keys on this instance, which is not common in the organization's environment. What did Marc do with these keys? Looking back at the AWS CloudTrail logs, the team sees that this access key spun up another instance, in a region the organization doesn't currently leverage. Was Marc trying to fly under the radar? Or did he accidentally script this instance creation and forget to set a region?

The final requirement from the internal audit organization is to explore what this instance had been communicating to. When the security team looks at the instance configuration further, it sees that the Amazon VPC Flow Logs show that this instance was communicating to a remote host over ICMP—an abnormal behavior. Fortunately, the team requires Amazon VPC Traffic Mirroring to be enabled on new Amazon EC2 instances that are created. This instance's traffic has been captured, so the team is able to analyze what was going over ICMP.

After further exploration, the team can piece together a timeline of events for its report to the requesting audit organization.

Summary

When moving to the cloud, it's best to outline a new incident response plan and plan out how you are able to perform investigations within AWS so that you can validate that any obligation you may have as a security organization can be met as well as it once was in-house.

With the fast and dynamic pace of the cloud, and with adoption of these new services increasing every day, security organizations need to review how they can adapt their processes and stay ahead to proactively enable developers and decrease risk in the environment.

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.

Sponsor

SANS would like to thank this paper's sponsor:



References

- ¹ US-Cert, "Computer Forensics," www.us-cert.gov/sites/default/files/publications/forensics.pdf
- 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

SANS 2019 Cloud Security Survey

Written by **Dave Shackleford**

May 2019

Sponsored by:

ExtraHop

Sophos

Sysdig

Executive Summary

Since our last Cloud Security Survey in 2017, we've seen a growing number of sensitive data disclosure scenarios and breaches involving the use of public cloud environments. One all-too-common scenario is sensitive data exposure in misconfigured and publicly available Amazon Simple Storage Service (S3) buckets. There are too many to name, but some of note include:¹

- A Verizon partner leaked personal records for more than 14 million Verizon customers, including names, addresses, account details and even account PINs in several cases.
- An Amazon S3 bucket leaked the personal details of more than 198 million American voters. The database contained information from three data mining companies known to be associated with the Republican Party.
- An ISP left 73GB of incredibly sensitive data in an exposed S3 bucket in late 2018 that included cleartext passwords, AWS keys, network diagrams and more.²

The Los Angeles Times exposed its website source code in S3, and in February of 2018, an attacker edited the code to include cryptocurrency mining functions.³ If the numbers are to be believed, 7% of S3 buckets are wide open to the world, and another 35% are not using encryption (which is built into the service).⁴ In June 2018, more than 22,000 container orchestration administration and API management consoles were discovered publicly, and some of them didn't have any authentication in place (and many had weak or default authentication in use).⁵ These primarily consisted of exposed Kubernetes platforms that security teams might not have had knowledge of or visibility into. Are these isolated incidents or common occurrences? What are security professionals doing to implement more effective controls within cloud environments?

The goal of the SANS 2019 Cloud Security Survey is to provide additional insight into how organizations are using the cloud today, what threats security teams are facing in the cloud, and what can be done to improve security posture in the cloud.

What stands out in 2019? Here are some of the key findings from this year:

- We saw a significant increase in unauthorized access by outsiders into cloud environments or to cloud assets; this occurred at 19% of organizations in the 2019 survey, whereas in 2017 this was experienced by only 12% of organizations.
- More than 55% of respondents in 2017 stated that they were frustrated trying to get low-level logs and system information for forensics, but only 30% said as much in 2019.
- ISO 27001 reports continue to be the most valuable audit reports made available by cloud providers, and more organizations are able to perform pen tests of their cloud provided environments than in the past.

What We're Doing in the Cloud

We asked the community what applications they have in the public cloud, and once again business apps and data top the list (76%). One big change we noted from our last survey was a significant decline in the use of workforce apps such as Dropbox. Only 45% said they were using such apps today versus the 84% who affirmed using such apps in 2017. This could be a simple difference in the respondents, given that SANS sees workforce apps as being a very popular category, so it's one to note and track for the future. Storage and archiving of data, as well as server (workload) virtualization in platform-as-a-service (PaaS) and infrastructure-as-a-service (IaaS) offerings, were also fairly popular. See Figure 1 on the next page for the breakdown.

About Our Respondents

This year, we had several hundred respondents who represent a number of industries. More than 21% are in the technology industry, and more than 11% each are in finance/banking and cybersecurity. Close to 10% are from government organizations, and many other verticals are represented in smaller numbers. Almost 40% work in smaller organizations (1,000 employees or fewer), more than 22% are in midsize organizations with between 2,000 and 10,000 employees, and close to 17% work in large organizations with 50,000 or more employees. Twenty-six percent of respondents are security analysts or admins, 12% are security architects, and 11% are IT managers or directors. Other roles represented include CSOs and CISOs, security managers and directors, and systems admins and compliance analysts. Organizations have operations in most countries, with the United States having the greatest presence (71%), followed by Europe (43%) and Asia (36%). Respondent organizations' headquarters are mostly in the US as well (62%), with Europe (18%) and Canada (6%) rounding out the top three.

The use of workforce apps, such as Dropbox, declined sharply since our last survey—with only 45% using such apps today versus 84% in 2017.

This year's survey also saw a consistent response in the number of public cloud providers that organizations are using. In both 2017 and 2019, the highest response category was "two to three providers." A higher percentage of respondents were using only one provider in 2017 (17%) versus today (16%). This slight change may indicate the beginning of a gradual shift toward multicloud. More organizations are using more than 20 cloud service providers in 2019 (7.5% total), versus our last survey, when just 4% used more than 20. See Figure 2 on the next page.

With the increase in use of cloud applications and multicloud implementations, particularly those that are oriented toward end users, we wanted to find out whether organizations are adopting new tools, such as cloud access security brokers (CASBs) and identity federation platforms, to help centralize control. Almost half of the respondents (48%) indicated they are using federated identity services to help centralize user access and authorization into cloud applications. Many are also using cloud network access services (43%) and CASBs (35%). Not as many organizations (19%) have adopted a multicloud broker to centralize access to PaaS, IaaS and other service provider environments. This makes sense. We need new services that can help centralize user access and identity, and also implement user-oriented policies for monitoring activity and protecting data (CASBs) as cloud application use grows.

As in past cloud security surveys, we looked at the kinds of sensitive data organizations are hosting in the cloud today. Business intelligence topped the list at slightly more than 48%, in a virtual tie with intellectual property (48%), and with customer personal information (43%) close behind. In 2017, business intelligence had come in second, behind employee records. This year, however, that former chart-topper had fallen to fifth place, with only 38% indicating that employee records are being stored in the cloud. Overall, the general trend is very similar to what we saw previously: Roughly one-half to one-third of organizations are willing to put a variety of sensitive data types in the




Figure 1. Cloud Applications in Use

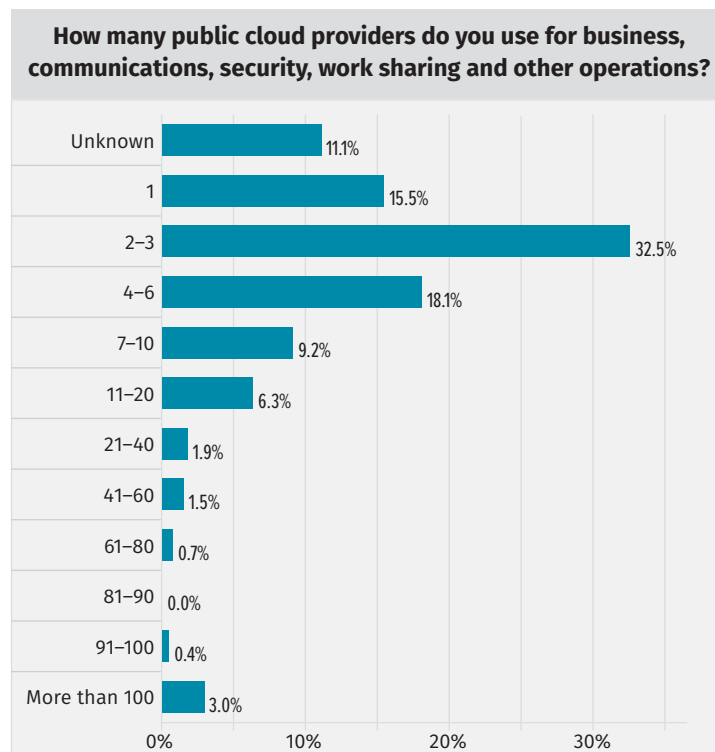


Figure 2. Number of Cloud Providers in Use

cloud, with lower percentages of some types (customer payment card information was less than 20% in both years, and health records were still lower than some other categories), as seen in Figure 3.

More than half of respondents (54%) indicated that privacy regulations such as the General Data Protection Regulation (GDPR) are impacting existing or planned cloud strategies, while 34% disagreed and 12% were unsure. Because of the GDPR requirements, organizations need to ensure cloud providers can adequately meet privacy compliance needs for some data types, especially consumer personal data. We cross-correlated those who answered yes to this question with the location of respondents' headquarters, and those who expressed the greatest concern were based in Africa, Europe and Latin/South America. This is not surprising, given that the European Union is directly affected by GDPR, and surrounding countries and business partners may be under pressure to provide the same protections.

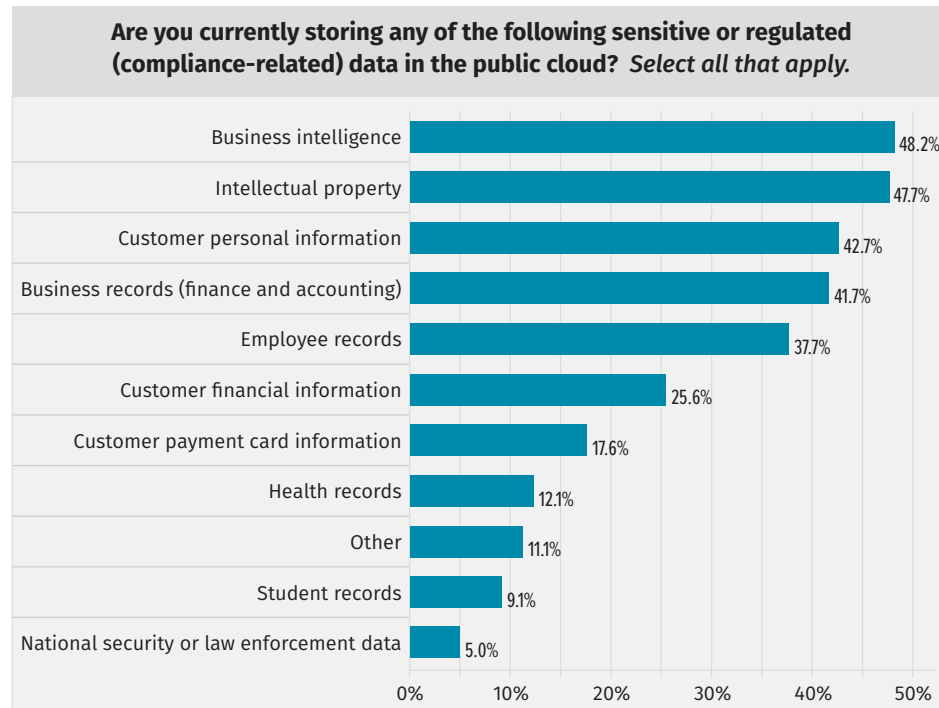


Figure 3. Sensitive Data in the Cloud

Concerns and Threats in the Cloud

As in 2017, unauthorized access to data by outsiders topped the list of concerns, at 56% (slightly lower than in 2017 but still the highest category). In second position, inability to respond to incidents (52%) moved up from seventh position in 2017, when 48% chose this concern. Other major concerns were lack of visibility into what data is being processed and where (51%, up from 48% in 2017) and unauthorized access to data from other cloud tenants, at 50% (very similar to our responses in 2017). The concern for data breaches by cloud provider personnel dropped from 53% in 2017 to 44% this year, which may indicate some growth in trust in the providers.

For the issues that were actually realized, downtime occurrences were fairly consistent from the last survey (up slightly from 18% to 21%). We also saw an increase in misconfiguration issues with application components and APIs. See Figure 4 for the full breakdown of concerns and actual incidents.

TAKEAWAY

The biggest change overall this year was a significant increase in unauthorized access by outsiders at 19%—in 2017 only 12% of respondents' organizations reported this problem.

More than likely, some of these issues go hand-in-hand. By exposing poorly configured applications and API interfaces (such as the Kubernetes APIs mentioned earlier), organizations are inviting access by attackers who are constantly using tools such as Shodan and network scans to look for targets. In 2017, the biggest issues that actually happened were downtime, misconfiguration and failure to meet service levels. While these are all still problems seen currently, they are overshadowed by actual attacks, which seem to have surged in the past few years.

Have these attacks and incidents actually led to cloud breaches in the past 12 months? Fortunately, the answer seems to be no for now—72% of respondents said they weren't aware of an actual breach, compared with 59% in 2017. This is good news, assuming that lack of awareness isn't an issue in itself. While 7% just aren't sure at all (compared with 21% in 2017), 11% said they did experience a breach, and another 11% think they've had one but can't prove it. The percentage of those who have (or believe they have) experienced a breach is roughly the same percentage as in 2017, which could be good or bad news, depending on how you want to see it: On one hand, things haven't gotten worse (superficially). On the other hand, why haven't we cut this number down in the past few years?

In 2017, we looked at what was involved in the successful attacks, and the top response was DDoS, followed by misconfiguration or other issues with hypervisors and virtualization management. The third major issue was the compromise or hijacking of credentials, but this was the No. 1 issue in 2019, with 49% experiencing this attack vector. Next in order was misconfiguration of cloud services or resources (42%), and then privileged user abuse (38%). These changes likely reflect the shifting

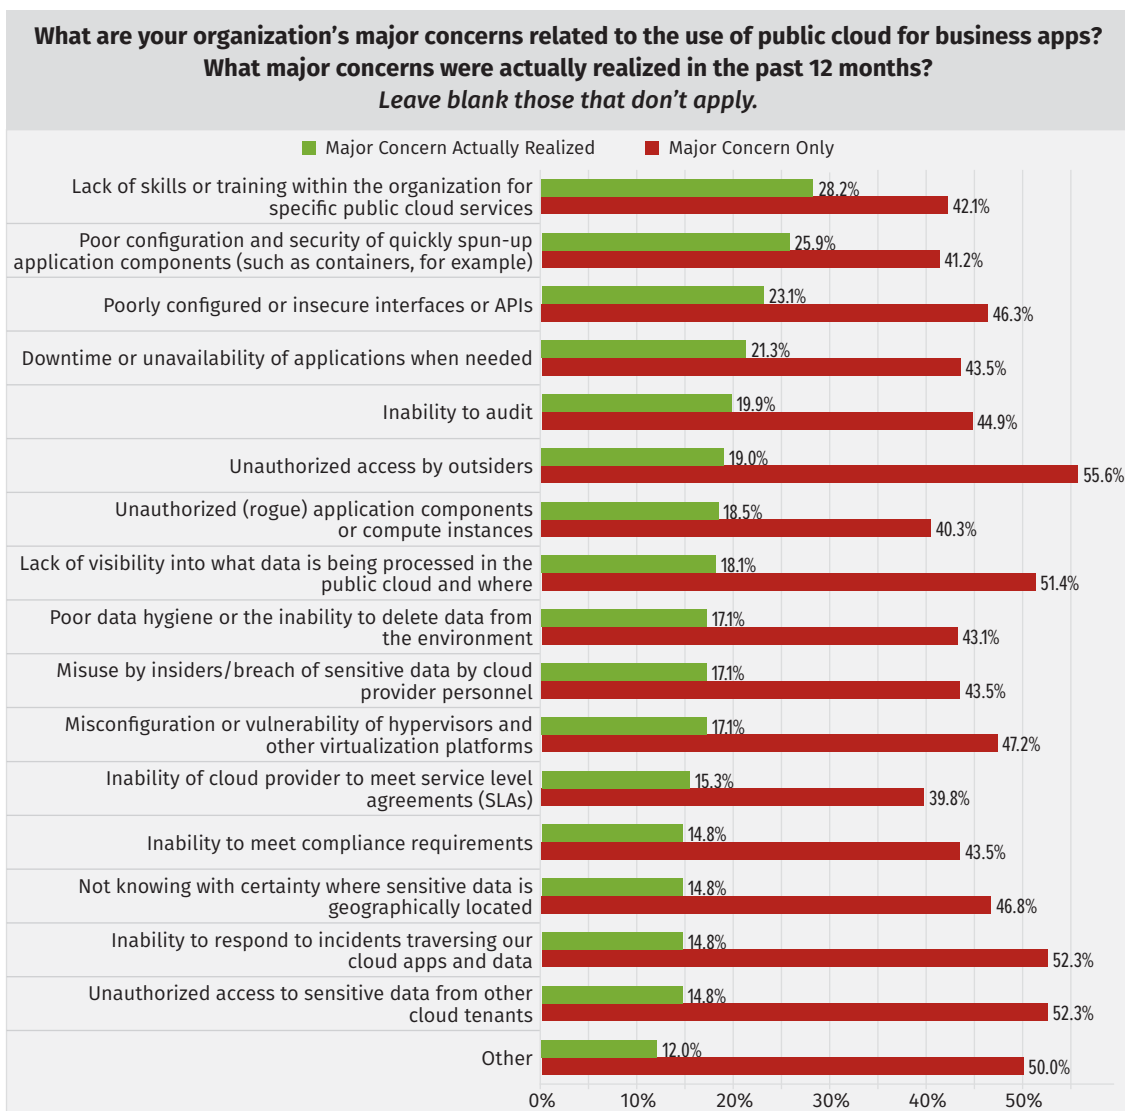


Figure 4. Concerns and Incidents in Cloud Today

nature of cloud, as well as maturity with providers and controls we have available to us. Virtualization elements are completely managed by public cloud providers, and so the surface area for attacks to this layer is greatly reduced.

DDoS attacks are still happening, but they don't seem as prevalent in breach scenarios due to improvements in DDoS protection from both the public cloud providers and the third-party services that have grown in popularity in the past several years. We're still not protecting credentials as well as we should, and misconfiguration of cloud resources is a pervasive issue, as evidenced by the plethora of exposed S3 buckets and APIs we see today. Privileged user abuse is likely symptomatic of the complexity of identity and access management (IAM) policies and settings that are tied to most cloud operations. The entire breakdown of things involved in the cloud attacks our respondents experienced is shown in Figure 5.

While it sounds as if most organizations haven't yet experienced breaches in the cloud, it may be too soon to know, given that many are unsure. This could also indicate a need for improved visibility into cloud and container environments overall. For those that did experience attacks or exposures, most of them related to credential hijacking and misconfiguration of cloud resources, which are both familiar issues to security teams.

Cloud Security Programs Today

As cloud use grows, organizations must develop and enhance their processes and governance models, so they evolve congruently. Today, 68% of organizations have cloud security and governance policies in place, which is up from 62% in 2017; 24% stated that they didn't, and 8% weren't sure. Gradually, we'll see more and more organizations evolve their governance and policy programs to incorporate cloud security and shared responsibility for controls and processes with cloud providers. In the types of attacks noted, only two would be wholly the responsibility of the provider: cloud provider vulnerabilities or API issues (20%) and hypervisor vulnerabilities or configuration issues (18%).




Figure 5. Cloud Attacks

Security Controls for Cloud Deployments

Through the years, we've seen teams get better at implementing some of the most common security controls for cloud deployments, but many types of controls are now available as security-as-a-service (SecaaS) offerings rather than standalone platforms. VPN was the most successfully implemented internally managed tool (59%), as it was in 2017. Network access controls and anti-malware were also touted in the 2017 survey as controls that organizations managed well internally, which again matches the results from this year (48% for network access controls and 50% for anti-malware).

In 2017, the top SecaaS controls in use were mostly the same, but anti-malware was used more frequently than network traffic analysis. Finally, the top controls managed between internal systems and SecaaS offerings in 2017 were vulnerability scanning and log/event management, where in 2019 the top results were log/event management and multifactor authentication. The full breakdown of controls in the cloud is shown in Figure 6.

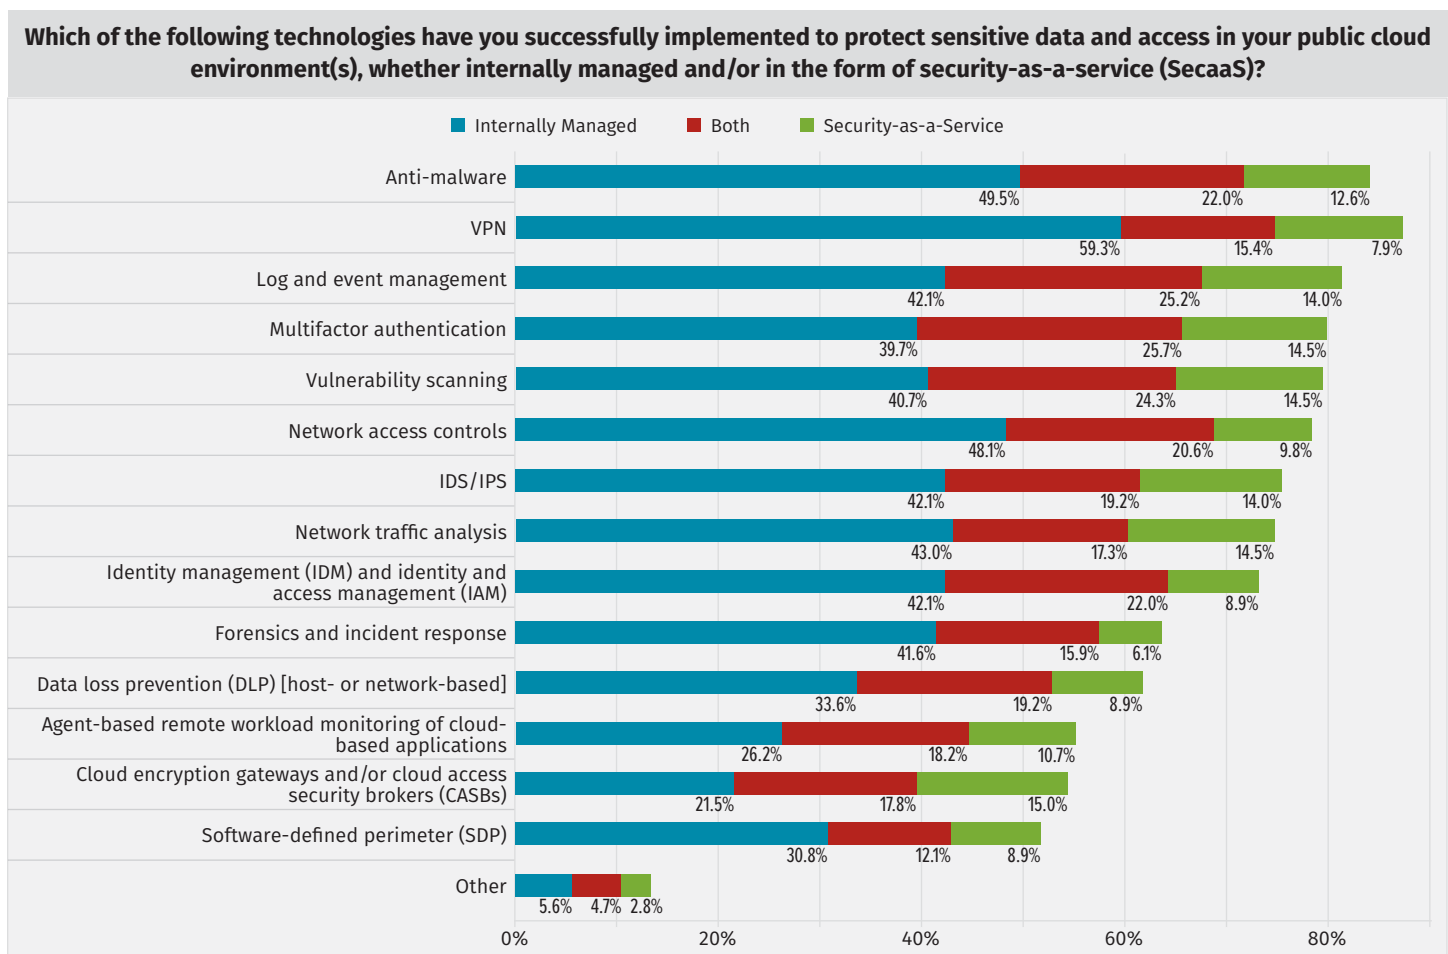


Figure 6. Security Controls for Cloud Adoption

There was a lot of interesting data with regard to controls. First, the majority of controls across the board are still being managed internally. In some categories, however, there has been more growth in a hybrid or services model, including CASBs and encryption gateways (18% for hybrid management) and identity management solutions (22% in hybrid management from slightly more than 16% in 2017). What stands out is the low numbers altogether. Many organizations may not feel wholly comfortable stating that these controls are capably implemented for the cloud yet.

This concern is somewhat substantiated by the fact that only 44% of respondents stated they are leveraging cloud provider APIs in the cloud to implement security controls (a critical element of automation and cloud security maturity)—almost unchanged from 2017 (43%). For those leveraging these APIs, the most common control is configuration management (75%), followed by logging and event management (72%), and then by identity and access management in third place (59%). These top three categories match what we saw in 2017, which suggests that these are the easiest to tackle through cloud provider-enabled API capabilities, the most critical for organizations to implement, or both. Collectively, though, all of these numbers are higher than they were in 2017, which is a positive trend; nonetheless, it is concerning to see fewer than half of organizations make use of the APIs provided. APIs offered by the cloud provider can afford security teams much more automated and capable access to and control over cloud environments, and hopefully we'll see increased use of these APIs in the future. See the full list of API-enabled security controls and functions in Figure 7.




Figure 7. API-Integrated Cloud Security Controls

Integration of Controls

Given that most organizations continue to manage many controls in-house, it's important to break down which controls organizations feel they've successfully integrated between traditional on-premises deployments and cloud environments, creating a true hybrid cloud security model. At present, 65% of organizations feel they've successfully integrated multifactor authentication, 58% feel that vulnerability scanning is well-integrated in a hybrid model, and 57% have anti-malware tools integrated. These findings are similar to the top three

technologies from our 2017 survey, though vulnerability scanning and anti-malware are reversed in order.

More than half have integrated network access controls (52%), and 47% have integrated network traffic analysis, which has been notoriously difficult in cloud provider environments in the past. Given many security teams' focus on capturing and analyzing network traffic for signs of intrusion and malicious activity, these are both critical to advance security maturity in the cloud.

Another 45% have integrated SIEM and event management tools, too. This is especially important, given that log and event management is one of the top three controls for cloud adoption (whether internally managed or through a SecaaS offering) and is a control area that involves high use of provider APIs. Because SIEM is a large, complex technology space, seeing its integration growing in a hybrid configuration is encouraging. The full breakdown of hybrid control integration is shown in Figure 8.

Note in Figure 8 that we also asked respondents which controls they planned to integrate in the next 12 months. Nearly a third indicated that they planned on integrating endpoint detection and response (EDR) tools (32%), followed by forensics and IR tools (28%), and then by event management at 26%. This indicates more focus on detection and incident response altogether, which has long been an immature control and process area for many teams.

In fact, we asked organizations what some of their biggest challenges were in adapting forensics and IR to the cloud. The top result was a lack of real-time visibility into events and communications involved in incidents—a problem that EDR and forensics/IR tool integration may help with significantly. Other major challenges cited include the difficulty in correlating events between on-premises and cloud environments (likely tying into the strong emphasis on SIEM and event management integration) and immature forensics and IR processes. Getting sound forensics evidence is also challenging, but it's interesting to note that in 2017, more than 55% of respondents stated that they were frustrated trying to get low-level logs and system information for forensics, and

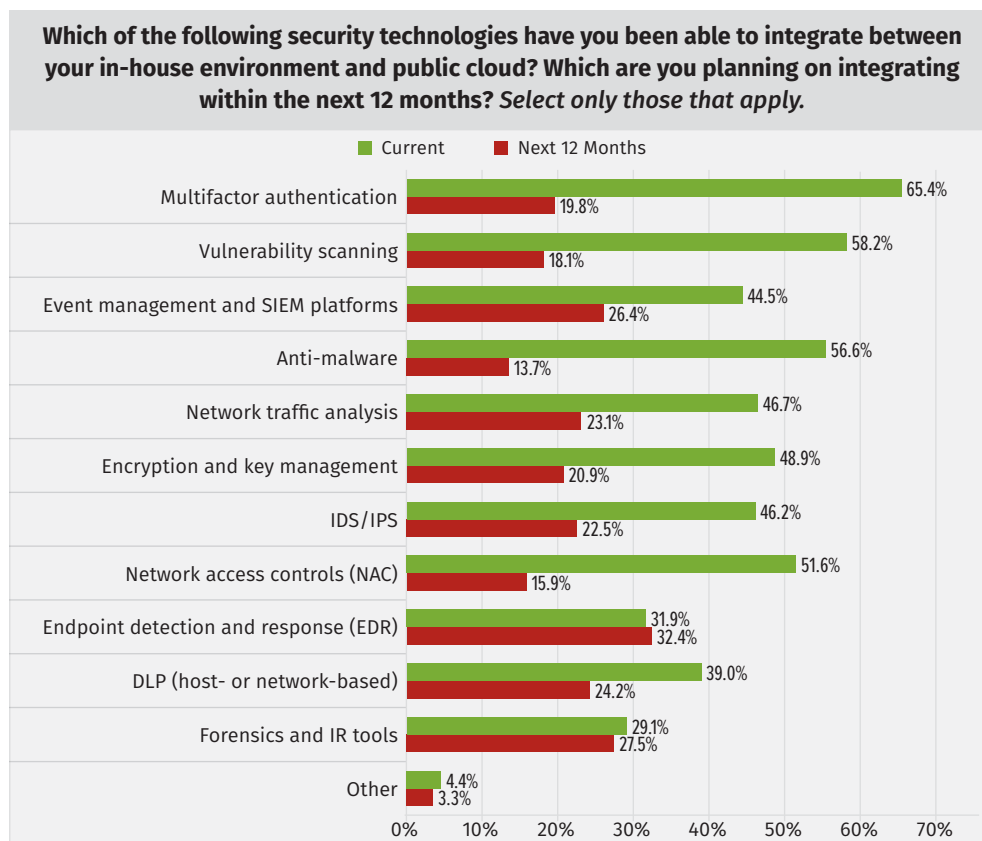


Figure 8. Hybrid Security Control Implementation

only 30% said as much in 2019. This is a strong indicator that providers are making this evidence more available than before, which bodes well for full integration of IR and forensics capabilities in a hybrid model in the near future. At its heart, this is a data security challenge as much as a visibility issue. The full list of forensics and IR challenges noted is shown in Figure 9 on the next page.

Returning to the concept of unifying and centralizing controls between on-premises and cloud environments, we looked to see whether security teams are finding any success in using the same vendors and technology providers across in-house and cloud environments for various controls. Unsurprisingly, respondents

provided the same types of answers mentioned earlier when expressing confidence in integrating control areas: multifactor authentication, network traffic analysis, vulnerability scanning and anti-malware. This is a strong indicator that success in implementing hybrid controls is likely linked to vendor products that integrate well in both environments, also providing central management capabilities. The same answers were given for plans to implement in the next 12 months, too (EDR tools and IR/forensics tools). See the full list in Figure 10.

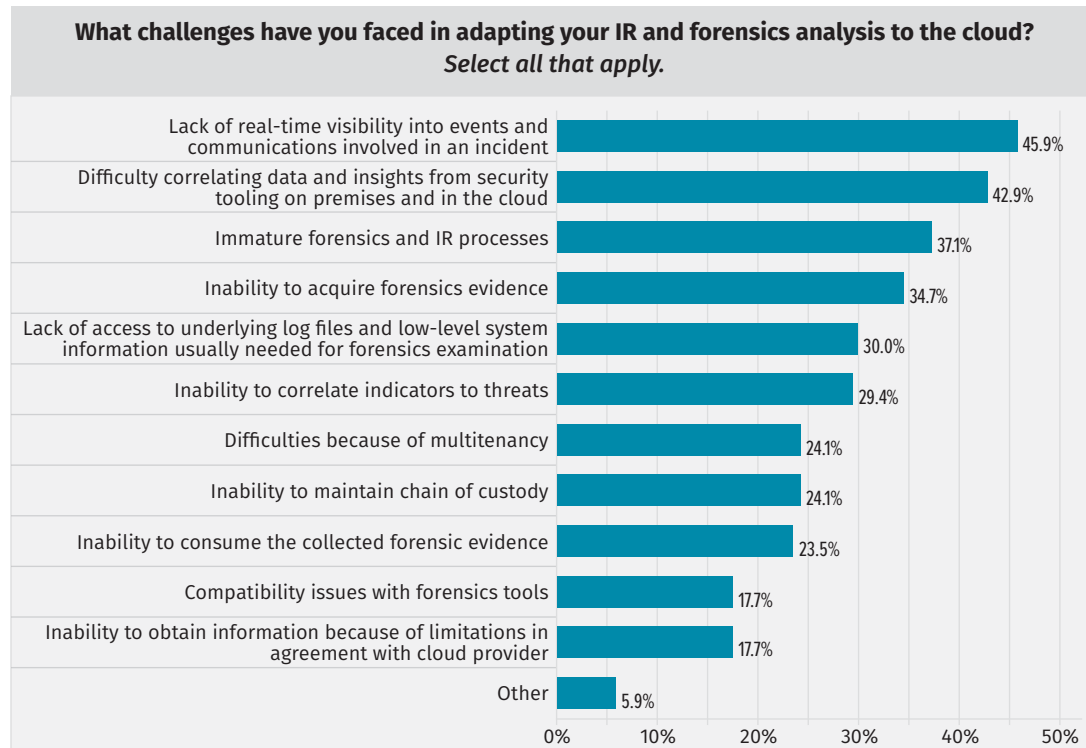


Figure 9. IR and Forensics Challenges in the Cloud




Figure 10. Single-Vendor Control Implementation for Cloud

Identity and Access Management

One of the most critical and growing areas of security controls for cloud environments today is identity and access management (IAM). IAM is rapidly becoming an essential element of most cloud implementations. More than half of respondents (52%) stated they were synchronizing in-house user directories to cloud-based directory services such as Azure Active Directory (Azure AD) and others, which is not surprising given cloud services' increasing reliance on access to user entities and attributes. On a related note, many organizations (35%) are also using identity-as-a-service (IDaaS) providers for SSO and federation activity to provision user accounts and attributes to numerous cloud services from a single source. More than a third of respondents (34%) use IAM policies to control object and application access and behavior, too—primarily in PaaS and IaaS clouds. Some are also mapping internal identities to their cloud providers and integrating traditional on-premises IAM suites to the cloud, as well, as seen in Figure 11.

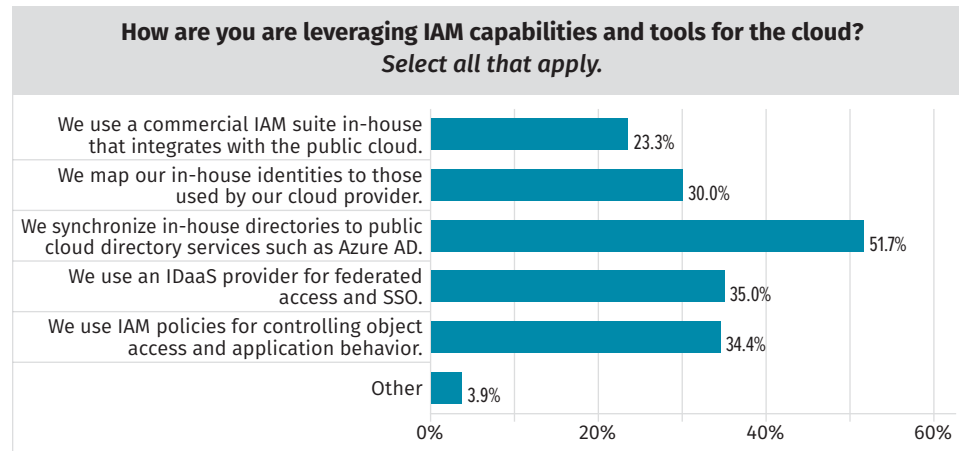


Figure 11. Use of IAM in Cloud

Automation and Orchestration

With a gradual shift toward dynamic asset creation and changes, as well as more DevOps-style application pipelines, security teams are seeing a definite need to implement some automated controls and monitoring tactics. A smaller subset of respondents (55%) voiced their thoughts on automation and integration tools and methodologies. Within that group, the most common tools in use today, selected by more the half of respondents, are template technologies for implementing infrastructure-as-code (AWS CloudFormation, Azure Resource Manager templates, Terraform, and so on). These allow security teams to build in cloud-native controls and monitor them as file contents, which can prove valuable in tracking and keeping up with highly volatile cloud environments. Security orchestration, automation and response (SOAR) tools are also in use by almost half of organizations, which presents a strong use case for central control and management of numerous security capabilities, ranging from detection to response. Configuration orchestration tools such as Ansible, Puppet and Chef are used by close to half of respondents as well, as are serverless technologies for execution of security functions. Not as many organizations have adopted security-specific plugins to build and deployment tools for DevOps

pipelines (Continuous Integration [CI]/Continuous Delivery [CD]). See Figure 12 for the full breakdown of automation/orchestration tools/methods in use today.

These are strong indicators that the use of automation and orchestration tools is growing, which is vital for security teams to keep pace with cloud operations and DevOps teams that want to move faster than ever before.

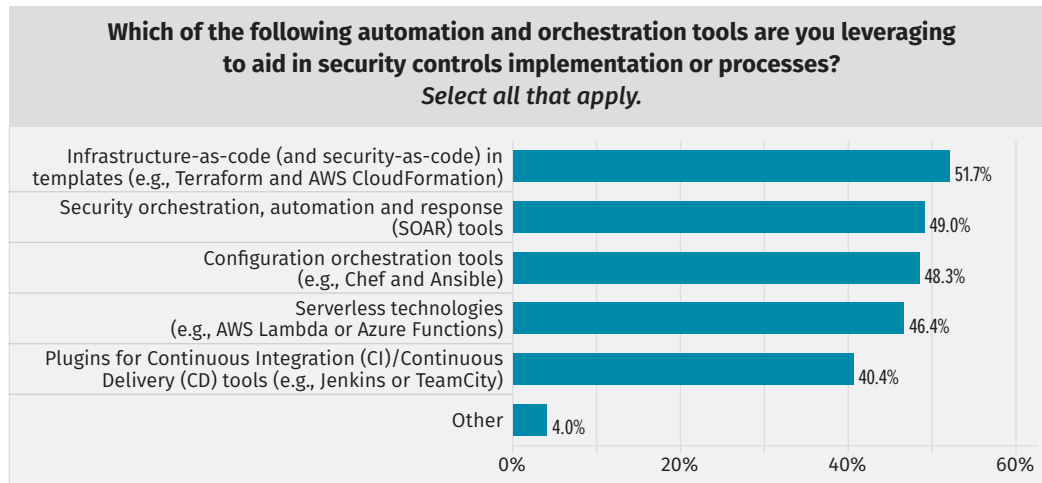


Figure 12. Security Automation and Orchestration Tools and Techniques for Cloud

Auditing and Assessing Providers

A consistent push in the security community has been to get cloud providers to document controls and provide more detail in the form of audit and attestation reports. We've consistently asked survey respondents to tell us which types of audit reports are most useful, because these are often among the few ways to assess what a provider is actually doing behind the scenes. Table 1 shows this year's results.

ISO 27001 was also the most valuable in 2017, but the CSA and SSAE reports were considered the second and third most valuable—the biggest change here is reporting reliance on the NIST Cyber Security Framework (CSF) and other controls, as well as FedRAMP for US government agencies and others to use a brokered, central auditing model in assessing and reviewing cloud provider controls. FedRAMP was considered valuable by only 28% of organizations in 2017, and has obviously grown significantly in maturity and adoption, likely due to increased adoption of the NIST standards in both public and private sector organizations.

Many organizations are also interested in performing penetration tests against their cloud applications and infrastructure. In fact, they might be required to do so for compliance reasons. Almost 55% of respondents stated that they are permitted to perform penetration tests against cloud assets (up from slightly less than 50% in 2017), while another 24% are not permitted to perform their own tests, but receive independent testing reports from the providers themselves. Only 10% are not permitted to test and do not get any reporting from the providers on pen test results (down from 18% in 2017, which is an

Audit and Security Reports	Percentage
ISO 27001	54.6%
NIST/FedRAMP	48.5%
SSAE 18 SOC 2	42.4%
CSA Cloud Controls Matrix and STAR program	31.5%
Others (CIS, PCI DSS, SIG, HIPAA)	6.7%

This year, nearly 55% of respondents stated that they are permitted to perform penetration tests against cloud assets, while just less than 50% of respondents had permission to do so in 2017.

improvement). Some types of SaaS providers do not allow pen tests because of the application environment configuration, but many PaaS and IaaS providers do. More providers overall are likely to facilitate pen tests in the future, to help clients meet internal standards or compliance requirements.

Conclusion

Every year, we conclude the survey by asking participants to provide general feedback on any other trends, concepts, experiences and issues they're seeing in the cloud today. This year, we also got feedback from the Cloud Security Alliance (CSA) as to what it is seeing in public cloud adoption and trends. Many organizations are continually evolving in their use of cloud services, looking to the cloud for procurement, management and other functions. The cloud provides capabilities for implementing new technology strategies in IoT and cryptocurrency, too, but many respondents mentioned the need for better APIs and automation capabilities to keep pace with the rapidly changing services offered. Especially as we shift toward multicloud deployments and cloud environments that are geographically dispersed, privacy issues are likely to become more of a concern. Many security teams aren't well versed in cloud concepts, both in design and operations areas and in DevOps/automation tools and tactics; this can be the case with container tools and technology, even more than with traditional server-oriented workloads. The perception remains that we aren't getting many needed details about security controls and capabilities from the providers, too, which limits our comfort level with the providers overall; conversely, some expressed the opinion that cloud may afford significant improvements in security over traditional on-premises data center environments.

Overall, the state of cloud security seems to be improving, albeit slowly. Cloud providers are becoming more open and accommodating of security data and controls, and more vendor solutions are able to bridge the gap between implementations on premises and in the cloud. There's progress, and more acceptance of in-cloud controls and services—but that progress is still slow.

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.

Sponsors

SANS would like to thank this survey's sponsors:



SOPHOS

Sysdig

References

- ¹ "Leaky Buckets: 10 Worst Amazon S3 Breaches," <https://businessinsights.bitdefender.com/worst-amazon-breaches>
- ² "Another S3 Bucket Leak—PocketiNet's Data Exposed!," <https://divvycloud.com/blog/s3-bucket-leak-pocketinets-data-exposed/>
- ³ "Guys, you're killing us! LA Times homicide site hacked to mine crypto-coins on netizens' PCs," www.theregister.co.uk/2018/02/22/la_times_amazon_aws_s3/
- ⁴ "7% of All Amazon S3 Servers Are Exposed, Explaining Recent Surge of Data Leaks," www.bleepingcomputer.com/news/security/7-percent-of-all-amazon-s3-servers-are-exposed-explaining-recent-surge-of-data-leaks/
- ⁵ "22K Open, Vulnerable Containers Found Exposed on the Net," <https://threatpost.com/22k-open-vulnerable-containers-found-exposed-on-the-net/132898/>