aws marketplace

SECURITY EBOOK

Six security trends to watch for in 2024

By SANS Senior Instructor
Dave Shackleford

Sponsored by AWS

In collaboration with:

SANS

Introduction

In the *2024 and beyond: Top six cloud security trends* webinar, Dave Shackleford, Senior Instructor and Analyst, SANS Technology Institute and Owner/Principal Consultant, Voodoo Security, talked about the hot topics that are dominating the cloud cybersecurity landscape. You can watch it here on demand.

In this ebook, Shackleford takes a deeper dive into the most important and prevalent trends in cloud security—including the rapid evolution of cloud-native application protection platform (CNAPP) and identity and access management (IAM) controls—that organizations should be tracking and implementing.

This ebook, sponsored by Amazon Web Services (AWS), also explains how you can build and deploy security solutions from independent software vendors in AWS Marketplace.

Browse AWS Marketplace to discover a variety of products that can enhance your overall cloud security posture.

Learn more by visiting AWS Marketplace ›

Plus, read on to the end of this ebook where you'll find more information about the AWS Partners and solutions that are mentioned throughout.

Dave Shackleford, Senior Instructor

SANS

Dave Shackleford is the owner and principal consultant of Voodoo Security and on the faculty at IANS Research. He has consulted with hundreds of organizations in the areas of security, regulatory compliance, and network architecture and engineering, and is a VMware vExpert with extensive experience designing and configuring secure virtualized infrastructures. Shackleford is a SANS Analyst, serves on the Board of Directors at the SANS Technology Institute, and helps lead the Atlanta chapter of the Cloud Security Alliance.

SANS | Research Program

Ebook

Six security trends to watch for in 2024

Written by **Dave Shackleford**

February 2024

aws marketplace

Introduction: What's Happening in Cloud Security

As the pace of cybersecurity threats and breaches accelerates, enterprise security teams struggle to manage risks and protect their environments against the evolving tactics and techniques attackers use to target cloud deployments. The cloud now brings a broader footprint of technology and assets into scope for security teams to discover, monitor, and protect. The constant and uncharted way in which workloads and cloud services are provisioned creates an expanding and dynamic attack surface that is hard to secure with tools and processes designed for legacy data centers. Security teams now have dynamic workloads with 10 to 100 times more containerized compute instances, large volumes of cloud assets with dynamic activity to track, and messy and overly permissive identity and access management (IAM) permissions to manage. Existing tools have not kept up with new tactics used by attackers in the cloud, leading to a weakened security posture. As a result, trying to develop, implement, and maintain a sound approach to cloud security has challenged many teams.

This rapid expansion of the potential attack surface has led to a large number of cloud vulnerabilities, misconfigurations, and security weaknesses. As cloud resources and sources of data increase, so does the burden of processing data into useful knowledge that can be applied to identifying, prioritizing, and remediating threats. Security, operations, and application teams are bombarded and overwhelmed by the number of alerts they face, leaving organizations with long exposure windows to critical vulnerabilities. Without real-time visibility across their entire cloud environment, it is difficult for teams to prioritize the most significant risks. It is critical for organizations to harden their security posture to prevent threats earlier in the application life cycle, especially as more organizations take advantage of DevOps practices common in the cloud, such as continuous integration and continuous deployment or delivery (CI/CD).

Bad actors are adapting to this new landscape and taking advantage of the growing vulnerabilities and security weaknesses. As more organizations are building new environments in the cloud or lifting and shifting to the cloud, the threat landscape has evolved to take advantage of these security gaps. In the cloud, attack patterns are different, with fewer traditional endpoint-focused attacks and many more attacks focused on the interconnectedness of software-based infrastructure, including identity assignment and orientation (both users and non-human, or machine, identities), application packages and libraries, exposed APIs, and more.

The overall landscape of cloud security today is increasingly complex, but there are new technologies and opportunities to improve security too. The evolution of cloud-native application protection platform (CNAPP) and IAM controls is happening rapidly, which should enable organizations to consolidate security monitoring and protection. Artificial intelligence (AI) and machine learning (ML) algorithms and tools are helping with threat intelligence, detection and response, automation, and much more. The philosophy of zero trust is actively manifesting in real-world technologies and capabilities that organizations can use, and this is nowhere truer than in the cloud.

Looking into 2024 and beyond, there's a lot to digest. This ebook will explore some of the most important and prevalent trends in cloud security that organizations should be tracking and implementing.

Cloud Risk and Governance as a Good, Combined Use Case

First, there will be a huge push toward a more coordinated governance and risk management strategy related to cloud in 2024. Organizations will continue to match information security requirements and the security capabilities of any cloud implementation to those of the mission/business needs being supported.

Security teams also will need to continually evaluate the threat profiles of cloud environments and the new types of issues that could occur. For example, S3 buckets being exposed is a big problem. How would you assess whether this affects you? Do you have S3 buckets? You may need different tools and processes in your assessment and review program to effectively handle newer cloud issues. Also consider how issues of logging, incident reporting, response, forensics, and other security-related functions should be addressed with respect to the service provider. In other words, match up your existing controls and requirements to those of the service provider in question and see where you may need new vendor products or capabilities that are offered natively. For example, could you use AWS Shield for distributed denial of service (DDoS) protection? Do you need to bring your Palo Alto firewalls into your AWS architecture? Are you going to get your cloud environment logs to your in-house SIEM, or do you need another product or service? This is really the longer-term process and maintenance you will need to keep up with to effectively look at the risks of any cloud project. Combining these precautions with a decent policy and a simple risk review process, you will be better positioned to maintain cloud security in an ongoing manner.

AI for Data Protection and Security Event Analytics

There are many potential benefits of ML and AI for risk management and security-oriented use cases today. Many AI risk management solutions rely on the mass computing scale achievable in the cloud, where large quantities of unstructured data can be analyzed and processed rapidly. In essence, risk management analytics that use AI can help organizations evaluate:

- Uncertain conditions or situations
- The likelihood of that condition or situation occurring based on context
- The effects the occurrence may have (the possible outcomes)

Additionally, risk management tools that make use of AI could be integrated into security automation workflows and could provide decision-making for security leadership during incidents, business continuity planning, fraud investigations, and more. Even with these benefits, there also are some potential drawbacks to using AI in risk management processes and practices. The first is cost. Processing large quantities of data, even using cloud-native services, can be very expensive, and specialized AI services also can be costly to enable. Second, many in the security community are concerned about data privacy with AI and ML. One reason for this is that the data organizations upload into cloud service environments may require data protection controls like encryption, transport security, tokenization, and obfuscation. Although most traditional data storage services in major cloud providers offer some or all of these, this changes significantly with specialized AI and ML services like AWS SageMaker and Amazon Rekognition (which uses AI to extract and analyze images and video), or Amazon Lex (AI for chatbots) and Amazon Comprehend for natural language processing. Not all these services can leverage the existing encryption key management and usage models and controls organizations may have deployed, so the data may be at risk of exposure. Aside from services in use, the geographic location of sensitive data used in ML and AI operations also is a major regulatory and compliance focus.

There are many different use cases where AI can benefit risk management and mitigation processes and practices. Here are some of the most common use cases today:

- **Threat intelligence analysis—**Threat intelligence data provides perspective on things like attacker sources, indicators of compromise, behavioral trends related to cloud account use, and attacks against various types of cloud services. Threat intelligence feeds can be aggregated, analyzed at scale using ML engines in the cloud, and processed for likelihood/predictability models. With attacks such as account hijacking and ransomware infections escalating, more rapid analysis of data and predictive intelligence could prove invaluable to security teams.

- **Security event management—**Log data and other events are being produced in enormous quantities, and security teams need to recognize specific indicators quickly, see patterns of events occurring, and spot events happening in the cloud environments where the events are occurring. ML and AI could easily augment massive event data processing technology to build more intelligence detection and alerting tactics.
- **Fraud detection—**For financial services firms and insurers, fraud detection requires an enormous number of inputs and data types and many intensive types of processing. Text mining, database searches, social network analysis, and anomaly detection are coupled with predictive models at scale, and cloud AI and ML engines likely could be an enormous help. This could be extended to things like fraudulent use of cloud services—for example, a cloud-based phishing attack from a hijacked account.
- **Employee workforce risk reduction—**AI and ML models can be used to process and analyze data related to workforce activities in high-risk environments like manufacturing plants, where accidents can prove dangerous or even fatal. AI algorithms could evaluate behavioral patterns noted before accidents occur, and also perform predictive scenarios to improve safety procedures and prevent incidents.
- **Data classification and monitoring—**Based on known content types and patterns, AI-based cloud analysis engines can process all data uploaded and created in the cloud environment to classify and tag it based on predefined policies, then monitor for access. Amazon Macie is an example of a service that uses AI for this purpose.

As the use of cloud-based AI and ML services becomes more commonplace, risk management teams undoubtedly will continue to benefit from the rapid analytics processing of large data sets, removing many limitations of the past's more manual risk management and risk analysis processes.

CNAPP: A Comprehensive Approach to Workload and Pipeline Protection

A cloud-native application protection platform (CNAPP), in short, is a combination of existing technology areas in cloud security. In general, CNAPP represents a convergence of workload security and configuration security for the cloud control plane, which cloud workload protection platforms (CWPP) and cloud security posture management (CSPM) already cover. To make things a bit more interesting, CNAPP also incorporates a bit of identity entitlement management, automation and orchestration security (particularly for Kubernetes), and API discovery and protection.

In many ways, CNAPP is focused primarily on the concept of "cloud native"—using cloud-centric technologies and controls that help lock down and secure the entire application deployment in a single solution. What's largely driving this in the industry is capacity. Most cloud security and security operations teams are already sorely taxed. They don't have the time or bandwidth to build and manage unique controls models that cover workloads, cloud services, identities, and the cloud control plane.

Many security, DevOps, and cloud engineering teams are asking why they may need a CNAPP solution. Is this something that makes sense? Are there any viable offerings in the marketplace? These are good questions, especially because we've seen a proliferation of cloud security tools and services emerge over the past several years. CNAPP is becoming a mature option in 2024. The component elements of CNAPP, described earlier, are rapidly maturing, but the combined offering in one solution is still in development. Today, most commercial solutions are good at one or perhaps several of the core elements that comprise CNAPP, but almost no commercial vendors excel in all of them. CWPP solutions are plentiful, CSPM services are common, but few vendors have really excelled in both these areas, along with orchestration and API security. Along with these capabilities, though, CNAPP is evolving to include:

- **Threat intelligence—**With broad insight into cloud-centric attacks against virtual machines, containers, and serverless workloads, CNAPP tools and services should provide insight into the types of attacks occurring in the cloud today.
- **Identity and privilege management—**As most DevOps pipeline tools, as well as cloud assets and services, have identity associations, CNAPP tools can help detect a variety of potential or existing security concerns related to privilege and role assignment and allocation.
- **Secrets management—**Cloud deployments can lead to an increase in services, tools, and platforms in use for DevOps and cloud engineering teams. This can also lead to increased use of privileged accounts, credentials, and keys that need to be managed and protected. In some cases, CNAPP tools can monitor and help protect these secrets and credentials.

In some ways, CNAPP is starting to emphasize cloud security controls and assessment earlier in the pipeline for more functions than standalone solutions have in the past. For example, CNAPP scans infrastructure-as-code (IaC) templates to look for configuration controls before deployment, and also looks for container image vulnerabilities and Kubernetes pod and cluster configuration settings. Although these areas all have been covered to some degree, no single vendor or solution has offered significant strength in all areas. CNAPP is also heavily focused on automation and API integration, which is appealing to DevOps teams who want security controls to be integrated with pipeline tools and services to minimize disruption and streamline CI/CD deployments.

CNAPP, as a concept, likely will succeed. There is a need for unified security capabilities across the DevOps pipeline (primarily in workload image, IaC, and orchestration configuration and vulnerability posture), configuration and controls for the cloud control plane, as well as runtime workloads in the cloud. Any solution in this space will need to offer powerful API integration for asset discovery and vulnerability/configuration posture and integration with DevOps pipeline tools to quickly and accurately assess IaC templates and workload images, and also be capable of runtime security protection for all types of workloads, including serverless functions.

AppSec in the Cloud: Protecting Against DDoS and API Attacks or Abuse

In the realm of cloud deployments, security teams are realizing that most assets and services are web applications. To that extent, we have to focus more on both APIs and availability attacks against exposed web application surfaces. With more and more integration and automation in cloud security becoming the norm, a new model for implementing security controls is emerging, and that model is entirely driven using security-oriented APIs.

Most cloud operations today can make use of available APIs, including configuration management and provisioning with tools like Ansible, Chef, and Puppet; development and code promotion into platforms like AWS CodeCommit; and security monitoring and management of identity data (via identity-as-a-service, or IDaaS, solutions); data protection and control with DLP and encryption; and network monitoring and traffic control. Most SaaS offerings make heavy use of APIs. In some cases, this may very well be the only way to integrate with SaaS and PaaS providers, like Salesforce and Office 365, to perform most functions related to data protection. For example, CASB and SASE/SSE providers like Netskope, Palo Alto, Zscaler and others can monitor data being sent to cloud provider environments and provide DLP and policy enforcement actions for usage of cloud services. Without deep API integration, none of this could occur.

Most large cloud providers offer native APIs to developers and operations teams alike—many with some related security functionality or advantages. One example would be Amazon Web Services' CloudTrail logging and event capture and retention service. Amazon has published a detailed set of APIs for CloudTrail[1] that allows security and operations teams to query events, list users associated with events, start and stop logging, and perform numerous additional functions. For security teams looking to build and maintain a DevSecOps workflow and supporting processes with development and operations teams, integrating into existing orchestration and automation design will require learning and understanding the APIs available and how they're being used.

[1] https://docs.aws.amazon.com/awscloudtrail/latest/APIReference/Welcome.html

Regardless of the provider and cloud services you use, it's important for security professionals to get acquainted with APIs exposed and available to developers and operations teams in the cloud. As development and operations tasks become more automated and integrated within cloud environments, the need to understand and leverage APIs will grow too. Of course, there are also potential risks associated with new APIs in cloud environments. In the Cloud Security Alliance's most recent Top Threats to Cloud research paper[2], one of the biggest risks cited is the exposure of insecure interfaces and APIs. Given that APIs are proliferating for provisioning, automation and orchestration, monitoring, and security functions, there is a need for security professionals to thoroughly assess and scrutinize these interfaces to ensure data and systems are not exposed or put at risk of compromise. Security professionals ideally need to scrutinize APIs for cloud providers as part of the procurement risk assessment process and push for contract language allowing penetration tests and vulnerability assessments against APIs offered by the providers. Emphasis should be placed on any API that handles sensitive data, and security teams should develop controls to ensure transport security, strong role-based access and authentication, message integrity and validation, and secure development practices whenever possible.

Many security professionals don't come from development backgrounds and may not be wholly comfortable with API analysis and security assessment. To help secure current and future cloud deployments, we'll need to get up to speed fast by engaging with development and operations teams, and also learning more about how APIs will be used in both cloud provider and brokering services. With traditional security controls changing or disappearing, there's no better time than now to get a handle on APIs and their role in tomorrow's cloud security architecture.

DDoS attacks often are initiated by malicious actors in an attempt to flood networks, systems, and applications with more traffic, connections, or requests than they can handle. Other types of DDoS attacks are more subtle, targeting specific services in ways that cause them to "hang" or fail. DDoS defense is a "must have" control for many organizations today. Attackers also are innovating and evolving DDoS attack strategies. As an example, a new zero-day bug dubbed "HTTP/2 Rapid Reset" was observed in a DDoS attack on August 28–29, 2023, that was apparently much larger and more impactful than any DDoS attack in history. Numerous large technology companies including Amazon Web Services, Cloudflare, and Google Cloud were targeted in the attack. They managed to minimize any significant outages through load balancing and edge security and availability controls, but the bug (CVE-2023-44487) has a CVSS score of 7.5 and could affect organizations that have not patched systems and applications. The bug is in the way HTTP/2 handles rapid requests and request cancellations in huge volume. By automating these alternating requests and cancellations, any exposed HTTP/2 services can be overwhelmed. For an example of scale, Google claimed that the attacks against them were reaching 398 million requests per second, several orders of magnitude larger than 2022's previous record DDoS of 71 requests per second.

[2] "Top Threats to Cloud Computing Pandemic Eleven," June 22, 2022, https://cloudsecurityalliance.org/artifacts/top-threats-to-cloud-computing-pandemic-eleven

There are many different types of DDoS detection and mitigation services, and a lot of variation between providers. When considering a cloud-based DDoS provider, look at:

- **Cost—**Some service providers charge a flat fee per month, with additional costs for mitigation and remediation. Others charge for quantity of bursting used, services employed, and number of sites/services covered.
- **Capabilities and features—**Organizations should look for DDoS defenses that encompass more than just volumetric attacks. Although these are still the most prevalent, many other types of attacks are emerging, particularly those targeting services and protocol behaviors, and any enterprise-class DDoS defense service should be able to handle some variations of all these attacks. In addition, many services will offer different management tools and reporting dashboards, as well as other features that may include threat intelligence briefings and additional research.
- **Response capabilities and experience—**Response SLAs are important, and so is experience. You should ensure that any DDoS response services are staffed by experienced analysts with a history of defending against large-scale and sophisticated DDoS attacks.

Cloud environments offer organizations a variety of strong DDoS detection and prevention capabilities, and services like AWS Shield Advanced also include advanced routing and mitigation techniques to help combat even severe floods of malicious traffic.

It is likely that DDoS attacks will continue in the future, whether fueled by criminal goals, political mischief, or other motives. At the same time, the severity and sophistication of these attacks is growing, and many organizations are not well-equipped to handle them. Using a cloud-based DDoS defense provider may prove to be an effective security control for preventing, detecting, and responding to these attacks, whether you have on-premises protection or not.

Zero Trust

Although there are many tools and controls available that can help monitor internal workloads and data moving between hybrid cloud environments, enterprises need to adopt one overarching theme when designing a dynamic security architecture model—one of "zero trust." To implement a zero trust model, security and operations teams need to focus on two key concepts:

1. Security will need to be integrated into the workloads themselves, and the actual behavior of the applications and services running on each system will need to be much better understood. The relationships between systems and applications will need more intense scrutiny than ever to facilitate a highly restricted, zero trust operations model.
2. End users will need behavioral monitoring and more thorough security controls in place to manage and mitigate account hijacking and other scenarios. Additionally, as hybrid cloud architectures become the new norm, many organizations are focusing heavily on automation, far beyond what we've traditionally seen in enterprise data centers.

By creating a layer of policy enforcement that travels with workloads wherever they go, organizations have a much stronger chance of protecting data anywhere the instance runs. In some ways, this shifts security policy and access control back to the individual instances rather than solely within the network itself, but hybrid cloud architecture designs don't easily accommodate traditional networking models of segmentation. Dynamic assets like virtual instances (running on virtualization infrastructure technology) and containers are difficult to position behind "fixed" network enforcement points, and so organizations can adopt a zero trust microsegmentation strategy that only allows traffic to flow between approved systems and connections, regardless of the environment they are in. Virtual systems can leverage a hypervisor backplane that all communications and behaviors are linked to, facilitating zero trust in a more scalable way. There are also physical models that accomplish this using specific network switches and connectivity platforms with built-in policy evaluation controls.

Zero trust microsegmentation prevents attackers from using unapproved connections to move laterally from a compromised application or system, regardless of environment. Essentially, zero trust facilitates the creation of "affinity policies," where systems have relationships and permitted applications and traffic, and any attempted communications are evaluated and compared against these policies to determine whether the actions should be permitted. This happens continuously, and effective zero trust control technology also will include some sort of ML capabilities to perform analytics processing of attempted behaviors, adapting dynamically over time to changes in the workloads and application environments.

By potentially eliminating lateral movement, a zero trust microsegmentation model also reduces the post-compromise risk when an attacker has illicitly gained access to an asset within a data center or cloud environment. Cloud design and operations teams (and often DevOps teams) refer to this as limiting the "blast radius" of an attack, as any damage is contained to the smallest possible surface area, and attackers are prevented from leveraging one compromised asset to access another. This works not only by controlling asset-to-asset communication, but also by evaluating the actual applications running and assessing what these applications are trying to do. For example, if an application workload (web services such as NGINX or Apache) is legitimately permitted to communicate with a database server, the attacker would have to compromise the system and then perfectly emulate the web services in trying to laterally move to the database server (even issuing traffic directly from the local binaries and services installed). These are just a few of the benefits of a zero trust segmentation strategy that can definitely help organizations implement granular access control policies across their internal and cloud data centers.

The other major model of zero trust we see emerging and evolving in the cloud is the security services edge (SSE). This model leverages cloud services providers as a core fabric for connectivity that end users rely on to get to both cloud-based and internal services. The core aspects of SSE include:

- **Zero trust network access (ZTNA)—**ZTNA focuses primarily on end user access to cloud and online services and data, with policies applied to evaluate who is accessing resources, from what system, and whether any behavioral aspects of access are suspicious or malicious. Key elements of ZTNA include:
  - Strong authentication and authorization of both endpoint systems and user accounts
  - Adaptive access policies that evaluate group membership and privileges, access behaviors, and known malicious/suspicious indicators
  - Browser isolation and sandboxing to prevent malware infection and other browser-based threats
- **Secure web gateway (SWG)—**SWG functionality includes content filtering and URL-based access controls, as well as some DNS monitoring and browser security controls. Most SWG platforms also include content monitoring and data loss prevention (DLP) policy tools. Leading solutions now offer remote browser isolation tools and capabilities as a means of providing a web browser "sandbox" for user access to designated sites.
- **Cloud access security broker (CASB)—**A CASB solution offers deep introspection to cloud services (primarily SaaS, but also applications and services in PaaS and IaaS environments) to look at API calls and behaviors to determine whether unusual activity is detected. Many cloud applications today are complex web services with vast arrays of API calls, and CASB services allow for much deeper analysis of specific interactions within the context of a single cloud application.
- Another capability some solutions tout is network traffic control, sometimes referred to as firewall-as-a-service (FWaaS), which replaces traditional next-gen firewall controls with a cloud-based model. This can be a valuable feature for SSE to control things like remote access protocols (SSH and RDP, for example), as well as any other non-web traffic that could be malicious.
- Today, most organizations need a suite of controls that can help implement a zero trust model of access control and monitoring, browser and cloud services security, and data protection for a remote workforce, and SSE provides this.

Identity and Access Management (IAM)

For organizations planning to build infrastructure in the AWS cloud, or to expand and enhance existing environments, AWS has developed a comprehensive model that can help inform design decisions known as the Well-Architected Framework. This framework consists of five pillars:

- Operational excellence
- Security
- Reliability
- Performance efficiency
- Cost optimization

Within the security pillar, one of the seven best practice design principles is *"Implement a strong identity foundation."* This principle includes several important themes in identity management, including implementing the principle of least privilege, enforcing separation of duties with appropriate authorization to resources, centralizing identity management where possible, and eradicating use of long-term static credentials across deployments. For any enterprise building infrastructure within the AWS cloud, all these concepts (and more) will be critical in designing and maintaining a sound security posture over time.

One of the primary tenets of the AWS Well-Architected Framework for IAM is to centralize identity and access wherever possible. This is a sound practice to pursue, for several reasons. First, when looking to manage provisioning and deprovisioning of identity accounts, it's ideal to perform account creation and revocation within a single, centralized identity store like Active Directory. This can help prevent local accounts from being created uniquely in cloud services, application stacks, or workloads, many of which may be forgotten or overlooked over time. Second, once a central identity repository is defined, it is ideally synchronized with a unified authentication and authorization portal (often referred to as single sign-on, or SSO) that can validate a user identity via credentials of various types and then facilitate access to additional resources from there. This approach is usually accomplished through federation standards for authorization, which most leading cloud services readily support. Another benefit of a centralized identity approach is reduced operational overhead and security governance.

All identities should dynamically acquire temporary credentials. For human identities, central SSO should be used to access AWS accounts. For third-party human identity access to your AWS resources, another service called Amazon Cognito offers lightweight "identity pools" that can be used to assign temporary, limited privilege credentials to access your AWS resources. For machine identities, organizations should rely on IAM roles instead of IAM users with long-term access keys. AWS Systems Manager is a more secure way to access and manage EC2 instances using keys or passwords, as instances leverage a pre-installed agent without any stored secrets present. The AWS Secrets Manager service also can be used to manage, rotate, and securely store encrypted secrets where short-term credentials aren't possible. IAM permissions can grant least-privilege access to secrets in Secrets Manager, and any API calls to access Secrets Manager secrets are logged in CloudTrail for auditing and monitoring.

For the past several years, cloud architects and security engineering professionals have been discussing the idea of "guardrails" in the cloud. Guardrails are usually implemented as defensive services and controls that are automated, continuously operational, and directly feed into detection and response processes and practices. Defining permissions guardrails is an important step in building a robust IAM architecture within AWS and it usually is focused on applying resource policies within IAM and assigning them to groups using centralized services like AWS Organizations. Guardrails are also important in enforcing a least-privilege model in the cloud, and there are several distinct types of identity-focused least-privilege orientations for cloud deployments and infrastructure. First, there should be a focus on any privileged users that need access to the cloud environment for administration, engineering, and security-focused tasks. Ideally, even in large organizations, this should be a relatively small number of users that are carefully set up and monitored.

Looking forward to 2024, IAM will definitely align and prove critical in developing zero trust use cases, and also ensuring privilege and secrets management is implemented through policy assignment in both third-party and cloud-native environments.

Conclusion: Where 2024 (and Onward) Is Headed

It's clear that, in 2024, we'll see all these trends continue to grow. Information security teams should be ready to discuss them with both internal and external stakeholders, particularly when sensitive data or potential risk exposure is involved.

In 2024 and beyond, a variety of trends are likely to grow and continue including:

- Improving cloud governance and risk management that aligns with more and more business objectives and use cases
- Convergence of tooling in cloud security beyond workload protection and configuration management, with maturation in the CNAPP space
- Major focus on IAM, especially centralized monitoring and control of identities and privileged identity control and oversight
- A trend toward zero trust within the cloud, aligning and focusing assets and workloads/applications based on a principle of least privilege and access minimization
- Cloud playing a major role in ML and AI, with data analytics and event management front and center

In all, these types of security controls and services are simply a natural evolution that reflects the nature of PaaS and IaaS software-defined cloud platforms and infrastructure. Security operations in large, distributed cloud environments will need to adapt to accommodate more dynamic deployments and changes, new services and workloads, and a significantly greater reliance on automation. In the next year and beyond, it's likely that all these trends will grow and mature significantly.

Sponsor

SANS would like to thank this ebook's sponsor:

aws marketplace

Start exploring

All of the featured solutions referenced in the ebook can be found in AWS Marketplace:

Cloud Risk and Governance

AWS Control Tower

WIZ

orca security

DRATA

JupiterOne

Thoropass™

AI for Data Protection and Security Event Analytics

Amazon GuardDuty

IBM

splunk>

sentra

DIG Digital Innovation Group

skyflow

CNAPP—Workload and Pipeline Protection

Amazon CodeWhisperer

CROWDSTRIKE

paloalto NETWORKS

aqua

sysdig

Strengthen your portfolio, predict risk, accelerate fraud detection, and augment advisory services all from a single destination, AWS Marketplace.

AppSec in the Cloud—Protecting Against DDoS and API Attacks or Abuse

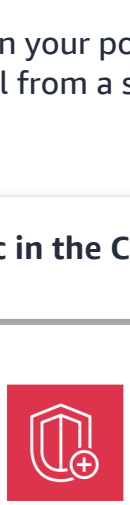

AWS Shield

F5

imperva

SALT

CEQUENCE

snyk

panther

Zero Trust

AWS Verified Access

zscaler

CHECK POINT

FORTINET

CATO NETWORKS

sonraí SECURITY

Identity and Access Management (IAM)

AWS Identity and Access Management (IAM)

okta

CYBERARK

TELEPORT

veza

Why use AWS Marketplace?

AWS Marketplace is a digital software catalog that makes it easy to find, try, buy, deploy, and manage software that runs on AWS. AWS Marketplace has a broad and deep selection of security solutions offered by hundreds of independent software vendors, spanning infrastructure security, logging and monitoring, identity and access control, data protection, and more.

Customers can launch pre-configured solutions in just a few clicks in both Amazon Machine Image (AMI) formats and SaaS subscriptions, with entitlement options such as hourly, monthly, annual, and multi-year contracts.

AWS Marketplace is supported by a global team of solutions architects, product specialists, and other experts to help IT teams connect with the tools and resources needed to streamline migration journeys to AWS.

How to get started with security solutions in AWS Marketplace

Security investigations and response teams use AWS native services and seller solutions in AWS Marketplace to help build automated, innovative, and secure solutions to address relevant use cases and further harden their cloud security footprint.

Here are some security solution resources to get you started:

Watch the webinar:

2024 and beyond: Top six cloud security trends

Watch on demand ›

Discover solutions:

Find security tools available to protect your AWS architecture.

Visit AWS Marketplace ›

Talk to an expert:

Speak with a solution architect who can help solve your business challenges.

Get connected ›