EBOOK

2023 Cloud Security Trends

A SANS ebook

Sponsored by AWS Marketplace

Introduction

In the past few years, major cloud providers are improving their security controls for customers and increasing visibility into their security solutions. Emerging trends such as endpoint detection and response and Zero Trust have focused the security community on the capabilities and services of both cloud and solution providers.

In this ebook, SANS Senior Instructor Dave Shackleford takes a deeper dive into six of the most significant trends to consider as part of your upcoming security strategy. Discover tools and better practices to help improve the security posture of cloud deployments.

You'll also discover how cloud marketplaces like AWS Marketplace help you find security partner solutions and expertise in these emerging trends.

Author: Dave Shackleford
Faculty – IANS Research
Owner and Principal Consultant – Voodoo Security
Analyst – SANS Technology Institute

Dave Shackleford, Senior Instructor

Dave Shackleford is the owner and principal consultant of Voodoo Security and 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. Dave 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.

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

SANS

Analyst Program

Whitepaper

2023 Cloud Security Trends

Written by **Dave Shackleford**

January 2023

aws marketplace

Introduction: The Changing Realm of Cloud Security in 2023 and Beyond

Over the course of the past several years, we've seen a number of major shifts in the realm of cloud security. Major cloud providers are improving security controls available to consumers and (maybe more importantly) providing more in-depth details about their internal security controls. In its 2019–2020 annual report on the top threats to cloud computing, the Cloud Security Alliance (CSA) noted several key changes from previous years' research. More importantly, the majority of concerns from the cloud security community focusing on the cloud providers themselves diminished dramatically. In other words, concerns about denial-of-service attacks against providers, virtualization security breaches and issues, and technology vulnerabilities in the provider environment were no longer considered top concerns at all, largely due to the rapid maturation of leading providers and the sense that the shared responsibility model is clearer and more well understood than ever before. This is great news for any organization moving to the cloud and/or increasing investments in cloud infrastructure deployments. Sadly, along with this news, the community noted some additional concerns that have grown in priority, namely the lack of proper oversight and attention on the part of cloud consumers to properly secure the cloud control plane itself.

This second factor has dominated the past few years, not only in capabilities and services offered by the providers themselves, but also in new focal areas and technology controls and capabilities from solution providers as well. There's been a significant increase in focus from the security community on cloud security, and several major trends have emerged that should carry us into 2023 and beyond. In the most recent SANS Cloud Security Survey,[1] we found that more organizations are moving critical applications, data, and workloads into the cloud. Organizations are more concerned than ever with unauthorized access to cloud resources by outsiders, poorly configured interfaces and cloud assets, and a lack of visibility into what's going on within the cloud. Cloud-centric attacks and breaches that occurred were almost entirely due to account compromise or poor configuration of cloud services and resources.

To better combat these challenges, the security community and the service providers integrating into cloud environments have worked to create better practices and tools that can help to improve the security posture of cloud deployments. This whitepaper describes some of the most salient and relevant trends we've noted in the past year.

[1] www.sans.org/white-papers/sans-2022-cloud-security-survey/

Trends in Core Cloud Infrastructure and Architecture

Many organizations are increasing the scope of cloud deployments steadily, and have been for some time. At the same time, large cloud service providers have expanded the catalogue of services and advanced cloud infrastructure and services, making it easier for a wide variety of IT and business teams to take advantage of cloud scale and capabilities. Today, more mature enterprises commonly make use of both platform-as-a-service (PaaS) and infrastructure-as-a-service (IaaS) capabilities in a unified cloud environment. PaaS services are customized system and application platform stacks offered to customers on demand. These usually take the form of standardized OS platforms or OS platforms with very specific application stacks installed, and they are often used for application development or customized hosting requirements. In almost every PaaS implementation, the provider offers a standardized set of APIs that are extended to customers. These are then built in to applications developed and deployed within the PaaS infrastructure.

With IaaS, a service provider hosts a complete virtual workload network for customers. With IaaS offerings, providers can support multiple distinct virtual machines (VMs) or interconnected VMs with supporting virtual infrastructure. In many environments, providers can also offer truly “private” networks and VMs for customers, where the consumer pays for a physically separate hypervisor server to avoid the security and performance risks inherent in multitenancy. Either way, in the IaaS model, consumers and providers share in the security responsibility. Consumers provision VMs and networks and must update and configure these components continuously. Providers support the hypervisor, underlying networks, and management components. A key takeaway is this: Consumers are responsible for all patching and configuration in IaaS environments, just like in their own data centers. In fact, many IaaS deployments resemble traditional hosting, but with the addition of dynamic provisioning and virtualization for scalability and rapid resource availability.

As this breadth of deployed services and assets has grown, cloud providers have steadily enhanced computing capabilities and services to make the cloud ecosystem even more tightly coupled, with a staggering array of innovative services and controls that can radically improve how IT infrastructure operates in a combined PaaS and IaaS cloud. With this expanded environment, however, we need to address new security requirements and responsibilities.

Maturing network detection and response in the cloud

Cloud environments should be developed with a robust defense-in-depth network security stack of controls, just as we've built in traditional data centers for many years. However, the types of controls will likely differ somewhat, including third-party systems and services, cloud-native controls, and cloud-focused network monitoring. In particular, we've seen significant growth in the realm of network detection and response for cloud-based environments, and that trend will definitely continue through 2023 and beyond.

Network detection and response (NDR) tools and platforms are network monitoring and traffic analysis engines at heart, focused on detection of suspicious and/or malicious traffic patterns that may indicate malware infection, active attacks, lateral movement between compromised assets, and more. NDR goes beyond traditional network intrusion detection (usually signature based) by analyzing patterns of behavior over a period of time in addition to specific network indicators such as IP addresses, host names, and ports/protocols observed.

Mature NDR platforms should ideally offer:

- **Network traffic monitoring capabilities—**NDR should be able to monitor raw network traffic in near real time (usually accomplished with traffic mirroring or taps to direct traffic to these platforms). Many NDR solutions can also monitor network flow data to develop behavioral patterns too.
- **Risk scoring and prioritization—**They should offer risk-prioritized views based on the confidence and severity of the detected network traffic and behaviors, as well as the business value of the assets affected. Mature NDR platforms can accommodate labels for critical network segments and networked systems/devices; for example, critical systems can be labeled as *crown jewels* to add weight to risk scoring for traffic to/from those assets.
- **Attack visualization—**Just like endpoint detection and response (EDR) platforms, NDR tools should provide click-down attack chain visualization tools to enable investigators to easily pivot on interesting data elements or drill down for more information. Whereas EDR will focus on local system events and indicators, NDR tools will show attack paths and communications observed as the major elements of visualization.
- **Network forensic collection capabilities—**NDR platforms should be capable of performing manual and automated packet capture and support for standard packet capture (PCAP) formats for additional analysis.
- **Deep analysis engine—**To detect network anomalies, leading NDR platforms should ideally include some machine learning capabilities and deep analytics processing to detect and correlate network behaviors. For many solutions today, this is done at least partially in a cloud environment.
- **Reporting and dashboards—**These should include severity and confidence indicators on threat alerts.

- **Response capabilities—**Some NDR solutions are capable of automated response efforts such as quarantine and blocking traffic, usually through partnerships and integration with firewall and other network traffic control solutions and vendors. Some vendors also offer endpoint agents that enable users to isolate systems.
- **Hunting/investigation tools—**To facilitate response efforts, NDR should offer security teams the ability to drill into traffic patterns observed and potentially recorded, search for particular patterns and content matches, and initiate investigations and open new tickets/cases.
- **(Optional) SSL decryption and inspection—**Some NDR platforms can terminate and inspect SSL traffic, whereas many others analyze traffic patterns and attributes instead of the content.

To facilitate sound network monitoring and traffic access, organizations should ensure they focus on two types of controls. The first type of network monitoring control organizations should enable within the IaaS cloud is collection of network flow data for monitoring communications to, from, and between workloads within virtual private clouds (VPCs). VPC flow logs enable monitoring and tracking of network events and behaviors on a large scale. With these types of flow logs, customers can designate a storage location for all logs to be sent and are also able to be aggregated and streamed to NDR services as needed.

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

Another capability many network security teams have sought in the cloud is full network packet capture controls. AWS VPC Traffic Mirroring enables the copying of network traffic from any compatible system in a VPC to a suitable endpoint such as elastic network interfaces (ENIs), network load balancers, etc. Many NDR tools and platforms can now leverage this mirroring capability to pull traffic from instances in AWS VPCs, allowing security operations teams to perform deep packet inspection, network forensics, and even selective packet filtering as well.

A variety of leading NDR providers integrate with AWS to offer robust network detection and response, including the following:

- ExtraHop
- Vectra.ai
- Gigamon
- Corelight Inc.

Tuning endpoint detection and response for cloud workloads

Today, we have a much better understanding of a sound stack of security controls and tools for securing cloud workloads. For more advanced endpoint security in 2023 and beyond, look at cloud-friendly and cloud-native options first, if possible. Many EDR vendors have adapted their agents to be very lightweight and supported in all cloud platforms, and these are good choices. Anti-malware technology should be selected from the cloud provider marketplaces and integrated with all images stored in the cloud provider environment. Cloud-focused endpoint posture management tools should also be integrated with all images to ensure they're active for any new workload deployments.

As you get started with locking down cloud workloads, keep the following in mind to help build a positive and productive feedback loop:

- Ensure that periodic reviews of the overall risk posture within cloud environments are performed to guarantee continued alignment of security and the other DevOps teams involved.
- Keep system instances in the cloud as locked down as you can, commensurate with the exposure and data classification types involved.
- Pay careful attention to privilege allocation and user, group, and role management associated with workloads. This can easily "creep" over time in a dynamic environment.
- Commit to a culture of continuous monitoring, helping to automate detection and scripted response activities that minimize manual intervention wherever possible.
- Discuss vulnerabilities detected in cloud deployments with all team members, and make sure DevOps teams are involved in vulnerability, patch, and configuration management discussions and policy creation.
- Discuss the changing threat landscape with DevOps teams, and solicit their feedback on practical measures that your organization can take to implement the most effective security without impeding progress or slowing down the pace of business activities.

For PaaS services such as containers and serverless, you need to incorporate controls that go beyond traditional runtime monitoring and management models. This should include native integration with the cloud provider environments and should be easy to automate for image check-in scans (and you definitely want to scan all new images immediately for vulnerabilities). For serverless functions, platforms that can be deployed for runtime monitoring and detection/response as separate serverless functions are usually desirable.

More modern EDR that is well integrated into cloud workloads should have capable policy controls that better reflect cloud attacks and security scenarios regardless. The following tools and policy controls should be available and mature for any cloud solutions considered:

- **Continuous threat monitoring and incident detection—**This is one of the most fundamental capabilities of any EDR platform, acting as the primary "eyes and ears" of workload monitoring. A wide variety of events can feed into this function, ranging from signature-based detection to anomalous behaviors observed.
- **Incident response—**Today's leading EDR solutions provide immediate response capabilities for cloud workloads. This is usually managed through command line and scripted queries or commands or through a central console from a solution provider. Automation techniques and integration with SOAR platforms are other key differentiators.
- **Threat intelligence collection/dissemination—**Security teams monitor a wide array of events occurring in the cloud environment, which range from workload events on Elastic Compute Cloud (EC2) instances, container events from Elastic Container Service (ECS), serverless events from Lambda or Fargate, and many others. Across all these events, teams can develop internal intelligence about actions of interest that can inform investigations and intelligence sharing with other internal teams or external parties alike.
- **Forensic evidence acquisition and analysis—**With significant visibility into what's happening across cloud workloads, security teams can assist in collecting and reviewing artifacts like logs, volume storage and/or memory snapshots, process indicators in memory from workloads, and more for forensics.
- **Threat hunting—**Threat hunting is the process of proactively looking for indicators, behaviors, and artifacts that could indicate activity worth investigating, often derived from internal/external threat intelligence. Security teams, with the workload visibility they gain from well-integrated EDR platforms, can help greatly with this security function.

Today, a variety of mature cloud-enabled solutions help to enable workload detection and response in the cloud:

- Aqua Security
- Trend Micro
- CrowdStrike
- SentinelOne
- Valtix

Cloud infrastructure growing as a critical element of Zero Trust architecture models and controls

In today's highly diverse technology environment, uniquely quantifying what *Zero Trust* means can prove challenging. For cloud-based scenarios, Zero Trust tends to be split between cloud-centric access control between and among resources and for user-based access to the cloud (for any variety of services). In the AWS cloud, the types of workloads, assets, and services have grown significantly, and numerous elements of access limitation between these cloud components are driving a renewed interest in Zero Trust architecture and controls. As a trend in 2023 and beyond, we're seeing the majority of Zero trust controls center on two areas. First, a variety of brokering services and client-based controls identify who/what the client is, where they're coming from, and what they're trying to access (and potentially when they're doing this) for access to AWS cloud services and applications (this is more focused on secure access service edge [SASE] and security service edge [SSE] types of providers). Second, for cloud services in AWS, a number of services and controls can monitor and manage lateral movement, unusual access attempts, and overly broad permissions and network access as well (the focus in this trend for 2023 and beyond).

To implement a Zero Trust model within the cloud, security and operations teams need to focus on two key concepts. First, security will need to be integrated into the workloads themselves. Second, the actual behavior of the applications and services running on each system will need to be much better understood, and the relationships between systems and applications will need more intense scrutiny than ever to facilitate a highly restricted, Zero Trust operations model. Organizations are also focusing heavily on automation (especially for discovery and monitoring), 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 regardless of where the instance runs. In some ways, this does shift security policy and access control back to the individual instances, versus solely within the network itself, but hybrid cloud architecture designs don't easily accommodate traditional networking models of segmentation.

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 includes some sort of machine learning 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 they contain any damage to the smallest possible surface area and prevent attackers 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) were 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).

Once the platform/tool of choice is selected, the next major planning element (besides installation) is policy design. Most of the leading providers of Zero Trust tools offer a form of "learning mode" that you can start out in, and that's definitely the right choice for almost all organizations—enable the Zero Trust engine and then monitor for what it sees. What you're looking to do is monitor which types of cloud applications and services are communicating between systems and network segments and map the communications to evaluate what is likely sanctioned and what might be malicious or unwanted traffic. When planning your policies, be sure to work closely with application and workload operations teams to better understand what is running in your cloud environment, as these teams will likely have a more accurate view of which communications should be in place. This way, you can build consensus on policy implementation before actually locking anything down.

At the same time, it's helpful to think about a "tagging" or "grouping" model that makes the most sense in your Zero Trust architecture. In other words, what systems are alike, and which systems should be communicating as part of defined application workloads? Common grouping strategies include business units (systems owned or maintained by a specific group or functioning as part of a business group), platform or application similarity (all databases or Windows servers, for example), and sensitivity levels (all systems in scope for PCI DSS compliance or those handling financial transactions). Choosing sound grouping for policies will enable organizations to implement them more quickly and effectively and may also make the policy design and governance discussions easier since you'll likely be working with existing teams that know how their applications should be functioning.

While cloud native services like Amazon EC2 Security Groups, network access control lists (ACLs), identity and access management (IAM) permissions, and others can help to establish a Zero Trust model within the cloud, highly capable solutions are also available to implement Zero Trust across workloads and cloud assets, including the following:

- Illumio
- Akamai
- Ermetic
- Drata

Identity and Access Management: Critical to All Cloud Deployments

Identity and access management (IAM) is really the practice of defining who needs access to what and then controlling the entire life cycle of user and access management across resources. For organizations planning to build infrastructure in the AWS cloud, or expanding and enhancing 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 different pillars:

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

Within the Security pillar, one of the seven best practice design principles is this: "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. Any enterprise building infrastructure within the AWS cloud will find all of these concepts (and more) critical in designing and maintaining a sound security posture over time.

Maturing cloud identity and access strategies

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 a number of 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 to 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, this is then 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 support readily. Another benefit of a centralized identity approach is reduced operational overhead and security governance.

For the cloud, organizations should carefully define how to approach the creation and life cycle of identities and groups. More progressive organizations have been building and implementing centralized cloud IAM teams to focus explicitly on this area of cloud security. For most organizations, it makes sense to start with an existing internal IAM team, if one exists. These teams often focus primarily on directory services like Active Directory, federation and SSO, as well as provisioning and deprovisioning users. These are all critical elements of a cloud IAM strategy, but additional IAM expertise and skills will be needed in adapting operating system privileges and permissions to cloud-based images and deployments, as well as configuring and managing cloud provider policy syntax and roles. To help in defining identities, permissions, and identity life cycles, organizations should build these teams from existing internal groups that already understand the business and goals of the organization, but sometimes cloud-specific IAM expertise requires recruiting from outside the organization. In addition to life-cycle definitions, defining naming conventions for machine identities is a smart endeavor to undertake, because unique naming identifiers can help to quickly discover these identities and monitor activity related to their use.

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 as well. For machine identities, organizations should rely on IAM roles rather than 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 can also 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.

While there are a wide range of cloud provider services and controls available to help define, implement, and manage permissions across all types of identities, many enterprises may choose to also leverage third-party solutions for optimized identity management and operations. Through marketplace offerings, several leading providers offer robust, centralized tools and services that can integrate with cloud SSO services like AWS IAM Identity Center and often support all leading federation authorization standards such as Security Assertion Markup Language (SAML) and Open Authorization (OAuth) for user and group mapping. Support for newer, flexible standards like the System for Cross-domain Identity Management (SCIM) v2.0 can help in automating the exchange of user identity information between environments too. Tools that can help to manage predefined permissions at scale such as AWS permission sets in AWS IAM Identity Center are also well suited to large deployment with a complex variety of identity use cases and stakeholders.

Continuously assessing and securing permissions within a cloud environment is important for ongoing security management. It is very common in the early stages of cloud deployments for application and DevOps teams to grant broad permissions and access to enable services and pipeline elements. Organizations need to restrict these to a least privilege model as soon as possible, and they should assess any new and updated permissions continuously to ensure that exposure is minimized. Access should be restricted to only assets, identities, and resources needed, and all permissions should be programmatically assigned and limited through centralized controls and services like AWS Organizations and other third-party solutions. Other guardrail services like AWS GuardDuty and IAM Access Analyzer can help to identify excessive or potentially dangerous permissions and IAM policy assignment as well. CloudTrail events can prove useful too, and IAM logs and analysis scans that AWS performs automatically in the background of every account provide access key and role last used information, as well as "last accessed" timestamps to identify unused users and roles and remove them.

In the AWS Marketplace, there are a variety of services that can help to centralize and manage identities and associated access models in the cloud, including:

- Okta
- OneLogin
- SailPoint
- Teleport

Improving privileged user management in the cloud

Defining roles, enabling strict access models, and limiting the resources available to users and systems are critical steps in enabling a sound cloud security strategy overall. A key element of IAM that security teams need to adapt to is the use of IAM for enveloping assets, allowing us to create "least privilege" architectures with affinity policies in place.

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

Several distinct types of identity-focused privileged user orientation for cloud deployments and infrastructure exist. First, there should be a focus on any privileged users who 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 who are carefully set up and monitored. The best practice for these users is to federate their internal user accounts directly to an assigned role within the cloud environment that has the fewest privileges assigned. The second major type of least privilege access model that all organizations need to consider is associated with deployment pipelines and associated systems and services. Whether on-premises or fully hosted within the cloud environment, deployment pipelines need certain privileges to update workload images and containers, access code repositories, assign metadata tags to resources, and monitor performance and security metrics and activities. The third major type of least privilege focus is mapping user, service, and application relationships wholly contained within the cloud environment. These might be EC2 workloads with instance profiles assigned that allow access to other AWS services like S3 buckets, Lambda functions that need to interact with CloudWatch logs and database services, or service IAM accounts/groups used to allow access between applications and services in the environment. Finally, security teams should carefully review privileges for accounts accessing other accounts' services when a multi-account strategy is in place.

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

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

As an isolation and segmentation technique, each account is a completely isolated set of resources that can be configured to access resources in other accounts. For multi-account strategies employed to limit the blast radius and provide highly granular least privilege access models, identity and access management is a critical element of managing the access between accounts too. AWS Organizations is a service that organizations can use to define policies to apply across multiple AWS accounts from a master control level. With Organizations, you can create service control policies (SCPs) that really govern the use of other IAM policies. Organizations can actually control the entire account, group, and role life cycle with regard to policy application, and can do so for accounts that need to interact or have some relationship. Some basic examples of how Organizations could be practical would-be governing business unit (BU) account use (because they may have totally different needs but still need some central control or billing) as well as governing and controlling DevOps and other team accounts (for the same reasons). Organizations is the lynchpin of a multi-account blast radius limitation strategy in AWS. Creating a centralized policy model within Organizations can allow security administrators to create different and least privilege policies for the appropriate accounts and assign them and/or revoke them easily. Organizations also provides a “master” rollup account that is often also the “payer” account that gets all “consolidated billing” for AWS accounts.

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 are usually focused on applying resource policies within IAM and assigning them to groups using centralized services such as AWS Organizations.

Many AWS Marketplace offerings can aid in centralizing and managing privileged accounts and identities in the cloud today:

- Sonrai Security
- Okta
- Ping

Cloud Threat and Vulnerability Management

In its 2019–2020 annual report on the top threats to cloud computing, the Cloud Security Alliance (CSA) noted several key changes from previous years' research. First, and likely most importantly, the majority of concerns from the cloud security community focusing on the cloud providers themselves diminished dramatically. In other words, concerns about denial-of-service attacks against providers, virtualization security breaches and issues, and technology vulnerabilities in the provider environment were no longer considered top concerns at all, largely due to the rapid maturation of leading providers and the sense that the shared-responsibility model is clearer and more well understood than ever before. This is great news for any organization moving to the cloud and/or increasing investments in cloud infrastructure deployments. Sadly, along with this news, the community noted some additional concerns that have grown in priority, namely the lack of proper oversight and attention on the part of cloud consumers to properly secure the cloud control plane itself.

Today, there are a wide variety of threats to cloud environments, with many traditional vulnerabilities (configuration errors, poor security in operating systems and applications) and exposure possibilities in cloud services themselves. Sadly, many of the issues are of our own making. Are we our own worst enemy? A trend we see growing in 2023 and beyond is a significant emphasis on cloud vulnerability and threat management. This will largely focus on attack surface management (ASM) and cloud security posture management (in other words, what's exposed and the condition of those assets).

Increases in focus on cloud attack surface management and posture management

Gartner defines external attack surface management (EASM) as "the processes, technology and professional services deployed to discover external-facing enterprise assets and systems that may present vulnerabilities." In many ways, EASM seems to be a new phase of maturation for the field of digital threat management, now including more vulnerability assessment and different types of cloud and other online assets. Leading use cases include:

- Exposed asset discovery and monitoring
- Attack surface exposure and risk validation
- Subsidiary and merger/acquisition risk (often related to cloud services and deployments)
- Supply chain risk tracking (also often tied to cloud services)
- Cloud-specific asset discovery and configuration management

This last category definitively overlaps with the more mature space of Cloud Security Posture Management (CSPM), tools that can assess the actual control plane of the cloud environments in use for compliance assessment, operational monitoring, DevOps integrations, risk identification, and risk visualization. A CSPM platform should continuously monitor cloud security risk and potentially implement configuration changes in the cloud environment that facilitate least privilege access and much more. These tools also offer threat detection, logging, and reports, and usually provide automation to address issues ranging from cloud service configurations to security settings as they relate to governance, compliance, and security for cloud resources.

Having interoperability between monitoring and automation is a critical advantage of a CSPM. For enterprises grappling with multi-cloud and container environments, where misconfiguration is a common threat to cloud security, a CSPM tool is an excellent step toward implementing continuous monitoring and alerting for the cloud provider fabric configuration, which will likely include identity controls, workload security, logging enablement, network configurations, and more. Organizations that are moving to or currently in multi-cloud deployment scenarios should strongly consider CSPM tools.

CSPM tools and services can monitor for a wide variety of issues within any cloud environment they monitor. The idea is to create a policy on what the “desired state” or “desired configuration” is for the cloud infrastructure and then monitor the actual state of what is in place. Examples of cloud control plane issues CSPM can look for include:

- No encryption enabled for cloud storage or databases
- No encryption for traffic in sensitive data in motion
- Lack of sound key management (old keys, stale keys, etc.)
- Poor IAM policies that don’t adhere to least privilege principles
- Privileged accounts without multifactor authentication enabled
- Open or permissive network access controls
- Exposed data storage, like accessible S3 buckets
- Minimal or no logging enabled within the cloud environment

When evaluating any sort of CSPM solution, security teams should look for key features that a mature service offering should provide:

- **Configurable and automatable remediation capabilities—**Ideally, any discovered issues can be remediated automatically or with minimal manual intervention.
- **Custom policy and rules engine enforceable across a multi-cloud environment—** The granularity and flexibility of a policy engine is one of the most important features for any CSPM solution. These need to properly and accurately assess cloud service provider settings and asset configuration.
- **Integration with DevOps pipeline stages and tools—**For any code or image repositories, build tools, etc., a CSPM platform should ideally be able to integrate and monitor activity here as well.
- **Detailed and configurable reporting—**Because CSPM is really a monitoring tool at heart, reporting is critical.

Applying CSPM to security operations should include the following:

- Asset inventory and classification (the faster and more accurate, the better)
- Focus on identifying access to the cloud control plane
- Monitoring policies for configuration and compliance
- Monitoring operational policies and configuration (performance, etc.)
- Collecting artifacts and insight into incidents for incident response (IR)
- Visualization and reporting of control plane risks

Whether vendors classify themselves as CSPM or EASM, it's going to be a growing area that many organizations focus on in 2023 and beyond.

Leading solution providers in the AWS Marketplace include:

- Palo Alto Prisma
- Wiz
- CyCognito
- Rapid7
- Aqua Security
- Immuta
- JupiterOne

Conclusion: Looking Ahead

As the types of cloud services available grow, and organizations continue to deploy large PaaS and IaaS environments that employ numerous interconnected services, the range of cloud security controls needed and surface to protect also gets larger. To keep up with the array of different cloud services in use, security teams need to learn and use more advanced controls and develop more dynamic and continuous processes for evaluating security conditions in their environments.

In 2023 and beyond, we see a variety of trends that will be likely to grow and continue:

- Cloud workload detection and response platforms that are more intuitive and tuned to cloud environments and potential attacks/threats
- Cloud network detection and response that takes advantage of packet mirroring and other strong access controls and monitoring available in large Paas/IaaS environments
- Major focus on identity and access management, 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 posture assessment tools for analyzing and remediating control plane security configurations and exposed asset vulnerabilities

In all, these types of security controls and services are simply a natural evolution that reflect the nature of PaaS and IaaS software-defined cloud platforms and infrastructure. Security operations in large, distributed cloud environments 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 of these trends will grow and mature significantly.

Sponsor

SANS would like to thank this paper's sponsor:

aws marketplace

AWS Marketplace

Cut software procurement time in half and realize an ROI as much as 550%

Why use AWS Marketplace?

AWS Marketplace is a curated digital catalog that simplifies software discovery, procurement, provisioning, and management.

With AWS Marketplace, customers can also utilize features that speed up production evaluation, improve governance and cost transparency, and enhance control over software spend. AWS Marketplace offers third-party solutions across software, data, and machine learning tools that enable builders to find, test, and deploy solutions to expedite innovation.

Explore and deploy solutions

IT decision-makers (ITDMs) cut their time in half using AWS Marketplace compared to other sources.

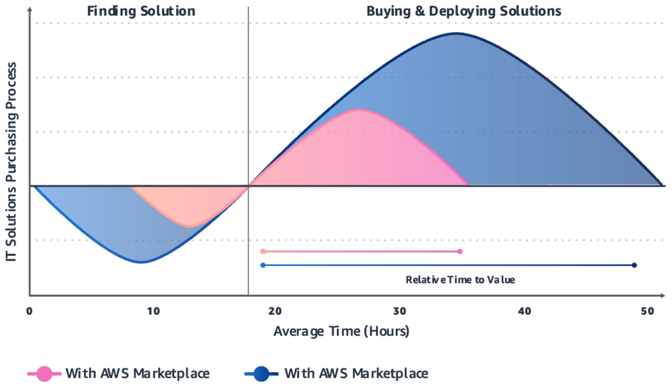

AWS Marketplace benefits

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.

Make more satisfying purchases

ITDMs feel 2.4x better about purchasing using AWS Marketplace compared to other sources.

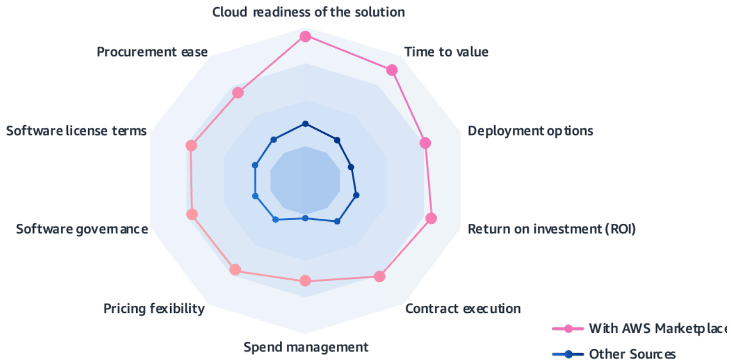

* Amazon Web Services (AWS) Marketplace surveyed 500 IT decision-makers (ITDMs) and influencers across the US to understand software usage, purchasing, consumption models, and compared savings.

Start exploring

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

Network detection and response

ExtraHop

Learn More

VECTRA

Learn More

Gigamon

Learn More

corelight

Learn More

Workload detection and response

aqua

Learn More

TREND MICRO

Learn More

CROWDSTRIKE

Learn More

SentinelOne

Learn More

VALTIX

Learn More

Zero Trust

illumio

Learn More

Akamai

Learn More

ermetic

Learn More

DRATA

Learn More

Identities and associated access management

okta

Learn More

onelogin

Learn More

SailPoint

Learn More

Teleport

Learn More

Privileged user accounts and identities management

sonraí SECURITY

Learn More

okta

Learn More

PingIdentity

Learn More

Cloud security posture management (CPSM) and external attack surface management (EASM)

paloalto NETWORKS

Learn More

WIZ

Learn More

CYCOGNITO

Learn More

RAPID7

Learn More

aqua

Learn More

IMMUTA

Learn More

JupiterOne

Learn More

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

The content and opinions in this ebook are those of the third-party author and AWS is not responsible for the content or accuracy of this ebook.