



Cloud Security for Amazon Web Services

The shift to cloud computing on Amazon Web Services (AWS) to modernize applications and infrastructure brings with it new requirements for security. Security teams must augment defense to stay a step ahead of fast-moving cloud threats. This guide highlights best practices for securing AWS environments and provides insights into how solutions from Sysdig and AWS help address key challenges, such as misconfigurations, supply chain risks, and evolving cyber threats.



Table of Contents

03

Introduction

Key challenges of securing the cloud

04

AWS cloud security approaches

Vulnerability management

07

Cloud security posture management

12

Cloud detection and response

15

Cloud security and generative AI

17

Conclusion

Introduction

In the cloud, every second counts. Organizations using Amazon Web Services (AWS) gain the ability to innovate faster, but must also ensure security keeps pace.

The speed of cloud-based attacks has surged as adversaries leverage automation and AI. In as little as 10 minutes, damage can be done. To stay secure, cloud security teams must find new ways to quickly and efficiently detect, investigate, and remediate threats faster than ever before.

Security practices are continuously evolving to address the unique challenges of the cloud. Point solutions have evolved to address a range of capabilities, from vulnerability and posture management to workload protection. At the same time, organizations are now finding value in adopting a platform approach to consolidate cloud security techniques to achieve a single view of risk across cloud, containers, and hosts.

Cloud-native application protection platforms (CNAPP) integrate a set of security capabilities that correlate signals to gain a more complete and efficient picture of your cloud and compliance from development to production. Unifying vulnerability management, posture management, permission management, threat detection, and incident response increases security efficiency and helps you stay ahead of cloud risk.

This guide offers a framework to help you navigate solutions for reducing risk and securing your AWS cloud investments. It offers insight into solutions available from AWS and highlights how the capabilities of Sysdig's CNAPP solution complement AWS cloud security.

Key challenges of securing the cloud

The cloud enables teams to configure infrastructure and deploy workloads with the click of a button. The pace of change opens the door to risk, especially as threat actors are also exploiting the speed of cloud automation to launch attacks in minutes. New and unforeseen visibility gaps can complicate security and compliance.

- **Misconfigurations and human error:** Misconfigurations are a significant concern in the cloud. Improperly configured cloud resources, permissions, and services can expose sensitive data or allow unauthorized access.
- **Software supply chain risk:** Adversaries target software during the development, distribution, or deployment process. Organizations often fail to implement effective security measures to prevent the introduction of malicious code or the exploitation of vulnerabilities.
- **Evolving cyber threats:** Cybercriminals are constantly evolving techniques to target cloud environments. Complex infrastructure with numerous entry points, including web applications, APIs, and user interfaces, can expose enterprises to risk if left improperly configured and unmonitored.

AWS cloud security approaches

Successful cloud security requires breadth of coverage across the software development lifecycle and depth of analysis to protect against known and unknown threats. Coverage from hosts and containers to serverless environments and the ability to correlate information with what's happening in real time with cloud services and identities is key.

Cloud security programs often emphasize two approaches: Shift Left and Shield Right.

- **Shift Left** approaches focus on processes and tooling that promote secure design and pre-release testing to identify security issues before they become production problems, and are intertwined with DevOps practices.
- **Shield Right** approaches focus on operational practices, security monitoring, and mechanisms to prevent security incidents, as well as detect and respond to events as they occur.

Shift Left and Shield Right security practices are both essential to securing your AWS cloud estate.

- **Cloud Workload Protection (CWP)**: Secure containers, Kubernetes, and hosts, identify and prioritize vulnerabilities, and protect AWS Fargate serverless workloads.
- **Cloud Security Posture Management (CSPM)**: Flag misconfigurations and automate their remediation. Continuously track security and compliance progress.
- **Cloud Detection and Response (CDR)**: Detect attack patterns across containers, Kubernetes, and cloud. Protect workloads against runtime threats.

The key practices incorporated in the above solutions all come together with a cloud-native application protection platform (CNAPP) for AWS. The following sections highlight the key capabilities of CNAPP — all of which contribute to helping you achieve end-to-end security in your AWS environment.

Vulnerability management

Vulnerability management is a crucial aspect of security for workloads running in the cloud. Considered key to both CSPM and CWPP practices, scanning for software flaws and known security issues is a must-have step in the application lifecycle to prevent security breaches.

New vulnerabilities are constantly being disclosed. Adopting a comprehensive vulnerability assessment approach is key to identifying and addressing issues throughout the software development life cycle (SDLC).

Assessing the security of modern applications requires finding and fixing problems early. This means scanning during development phases and continuing to scan for vulnerabilities through runtime. Checking for issues at each stage helps identify vulnerabilities missed during earlier stages, those introduced during runtime, as well as risks disclosed after your last scan.

Full lifecycle vulnerability management means scanning for issues at different stages and locations:

- Local scans on developer machines
- Continuous integration and continuous delivery (CI/CD) pipeline scans
- Registry scans
- Runtime scans

Shift-left security challenges

Shift-left security has a noise problem. Organizations start their journey with pre-release scanning tools and quickly drown in a deluge of scanner output. It's a struggle to find an efficient way to pass or fail application releases.

Development and security teams must sift through and vet findings to prioritize flaws that are severe and actionable. The work is tedious, stealing time from more important tasks. Addressing the pitfalls of security testing isn't simple. Teams need as much contextual information as possible to begin to reason about the risk associated with each finding.

AWS vulnerability management solutions

AWS provides solutions that identify software vulnerabilities to help cloud teams detect and route vulnerability findings to the appropriate teams to improve security and compliance.



Amazon Inspector detects software vulnerabilities and unintended network exposure in AWS workloads, such as Amazon EC2 and container images in Amazon ECR, and within continuous integration and continuous delivery (CI/CD) tools.



Amazon ECR offers basic scanning — either on push to a repository, or manually — using the Common Vulnerabilities and Exposures (CVEs) database from the open source Clair project. ECR also integrates with Amazon Inspector to provide enhanced scanning,

Vulnerability management for AWS with Sysdig

Sysdig incorporates vulnerability management into the various stages of the development process, as well as runtime. In addition, it helps AWS users prioritize risk and consolidate host and container image scanning into a single workflow to save time and money.

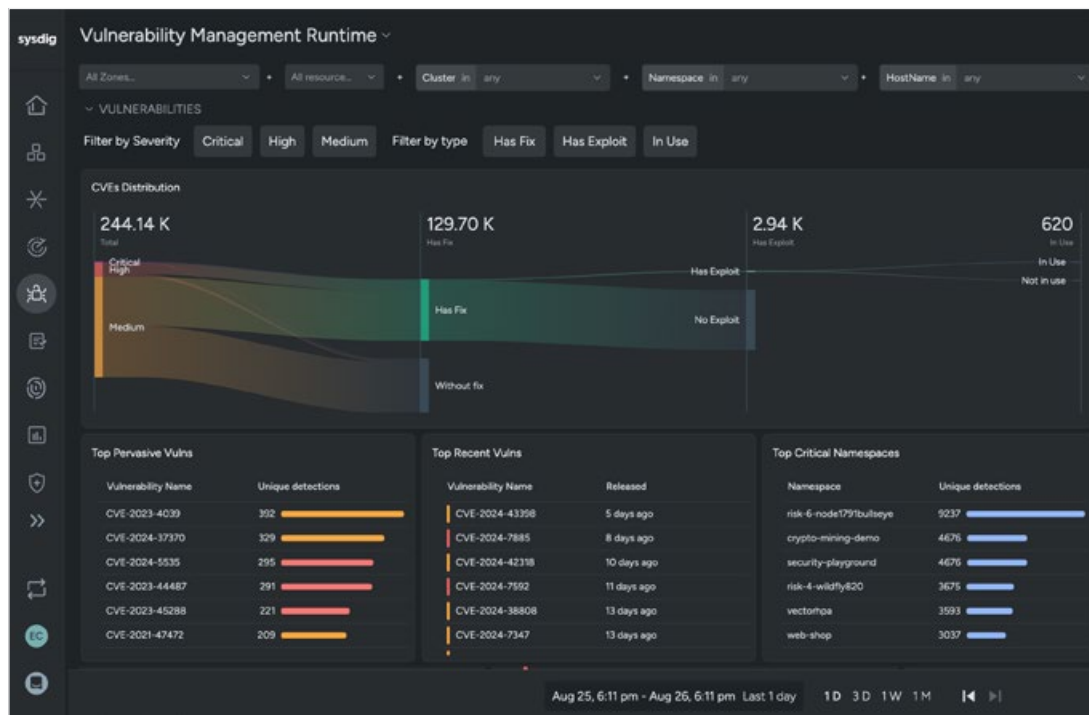
- **CI/CD pipeline scanning:** Vulnerability scanning in the CI/CD pipeline evaluates container images as a build step before pushing to a registry. As an additional gate, you can fail builds that don't pass the security policy evaluation.
- **Registry scanning:** Registry scans ensure that container images and artifacts are checked for vulnerabilities before running in production on AWS.
- **Runtime scanning:** Scans for vulnerabilities at runtime to identify issues not identified during earlier stages, introduced during runtime, or disclosed after the last scan occurred.

Prioritizing vulnerabilities with runtime insights

To help teams stay ahead of vulnerability threats, adding context to identify in-use vs. dormant vulnerable packages helps reduce noise and spotlight actual risk. Sysdig's runtime insights deliver this visibility by profiling containers and making this information available to vulnerability management and application security (AppSec) tools.

Runtime insights are accessible within Sysdig's CNAPP, but are also used by third-party AppSec solutions. Industry AppSec leaders, including [Snyk](#), [Checkmarx](#), and [Docker](#), have integrated with Sysdig to deliver the ability to prioritize in-use vulnerabilities and eliminate threats faster.

Figure 1:
Runtime Vulnerability
Management



Agentless and agent-based scanning

Organizations are often hesitant to install and maintain agents to deliver needed functionality. Where possible, teams want to utilize an agentless approach. Agentless security scanning generally relies on cloud provider APIs to collect information and perform vulnerability assessments.

The drawback of agentless scanning is that it typically lacks real-time visibility. This means that teams may lack information about intermediate states of the system between scans. In addition, agentless is typically unable to provide insights into whether a vulnerable package is in use in a running environment.

Sysdig integrates both agentless and agent-based options:

- Agentless scanning leverages AWS APIs to discover and scan resources.
- Agent-based scanning uses a lightweight package to provide both node-based scans and runtime visibility.

Both options can be used in tandem. With this approach, the agent profiles workloads to identify in-use packages. This information is used by the agentless scanner to prioritize vulnerabilities.

To learn more, read [Securing the Cloud: A Guide to Effective Vulnerability Management](#).

Cloud security posture management

According to Gartner®, “through 2025, over 99% of cloud breaches will have a root cause of a customer misconfiguration or mistake.¹” Cloud Security Posture Management (CSPM) helps provide visibility into your cloud configurations to identify and remediate risks and to proactively safeguard your AWS environment.

Ensuring secure cloud configurations

Posture management is a cornerstone of any cloud security strategy. Cloud misconfigurations leave your business exposed to risk. Wrongly configured hosts, container runtimes, clusters, storage, or cloud resources create an easy way to escalate privileges and perform lateral movement. Evaluating your AWS cloud accounts and services against benchmarks and posture controls helps you detect when resources deviate from security best practices.

Rather than performing manual evaluation and remediation of cloud configurations, CSPM solutions can automatically assess the state of your cloud configurations and provide a readout of risky misconfigurations. In some cases, CSPM will also automate remediation by updating or disabling flawed configurations.

Securing infrastructure as code (IaC)

Infrastructure-as-Code (IaC) using tools like Terraform and [AWS CloudFormation](#) has become a core component of IT provisioning and administration in the cloud. Validating IaC configurations is another key component of CSPM.

IaC security tools and practices allow engineers to find and remediate security problems within IaC templates. The goal is to minimize the risk of inadvertently introducing security problems via IaC. IaC security is a part of posture management in that it introduces governance designed to mitigate security risks.

Kubernetes security posture management

Kubernetes security posture management, or KSPM, is the use of security automation tools to discover and fix security and compliance issues within any component of Kubernetes. You can think of KSPM as CSPM for Kubernetes environments. KSPM analyzes Kubernetes resources and host configurations, in addition to Kubernetes audit logs, to help you prevent and remediate security risk in your cloud-native infrastructure.

With a cloud-managed Kubernetes service such as Amazon Elastic Kubernetes Service (EKS), AWS manages the Kubernetes control plane. Therefore, security posture is managed by AWS. For other aspects of an EKS environment, such as worker nodes, your organization is responsible for hardening, patching, and managing the security updates. For a fully self-hosted Kubernetes environment, you'll have to manage the security posture of the complete environment.

1 Gartner, Risk-Based Evaluations of Cloud Provider Security, Charlie Winckless, Jay Heiser, 16 January 2023.

GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved.

Managing permissions and entitlements

Overly permissioned cloud accounts and roles pose another critical security problem. Identity and Access Management (IAM) is critical for helping AWS users lock down access to avoid the risk of data breach, privilege escalation, and lateral movement.

As you increase the use of services and features on AWS, it becomes harder to know exactly what the least privileged entitlements should be. Permissions are often misconfigured, allowing unnecessary access rights.

Carefully assigning the correct permissions is fundamental to addressing identity risks in the cloud and achieving the practice of least privilege in your AWS environments. Cloud Infrastructure Entitlement Management (CIEM), considered a key CSPM capability, is specifically designed to address the complexity that arises with keeping up with permissions in cloud environments.

98%

Sysdig's 2024
Cloud-Native Security
and Usage Report
found that 98% of
permissions granted
are unused.

AWS CSPM, IaC, and CIEM solutions

AWS solutions for posture management and code security help cloud teams aggregate security findings, analyze configurations and code, and identify permission issues.



AWS Security Hub

performs security checks and aggregates security alerts by accounts and resources, helps cloud teams run security best practice checks, visualizes findings, and automates remediation.



Amazon CodeWhisperer

scans your code to identify security vulnerabilities and provides code suggestions to help configure best practices and remediate identified issues.



AWS IAM Access

Analyzer provides capabilities to set, verify, and refine your cloud permissions, and validates that your policies match your specified corporate security standards.

Posture management for AWS with Sysdig

Sysdig's CSPM solution continuously manages cloud infrastructure and identity risks by identifying and enabling the remediation of misconfigurations in the cloud control plane, cloud resources, cloud-deployed workloads, and permissions.

Security best practices and compliance

Sysdig helps you proactively assess target environments against security and compliance standards, common frameworks, regulatory requirements, and your internal company policies.

Built-in posture assessment tooling includes the following standards:

AWS Well-Architected Framework	AWS Foundational Security Best Practices	Center for Internet Security (CIS) Benchmarks	Defense Information Systems Administration (DISA) Security Technical Implementation Guide (STIG)	Digital Operational Resilience Act (DORA)
Federal Risk and Authorization Management Program (FedRAMP)	General Data Protection Regulation (GDPR)	Health Insurance Portability and Accountability Act (HIPAA)	Health Information Trust Common Security Framework (HITRUST CSF)	ISO/IEC 27001
National Institute of Standards and Technology (NIST)	Network and Information Security (NIS) Directive (NIS2)	NSA/CISA Kubernetes Hardening Guide	Payment Card Industry Data Security Standard (PCI DSS)	System and Organization Controls (SOC)

Figure 2:
AWS cloud security posture report

Result	Requirement / Control	Controls F...	Policy / Control T...	High	Med	Low	Accepted	Passing
✖	[S3.2] S3 buckets should prohi...	1/2	AWS Found...	4				
✖	S3 - Blocked Public Access (Bu...		Resource	4				
✔	S3 - Blocked Public Access (Ac...		Identity					
✖	[S3.3] S3 buckets should prohi...	1/2	AWS Found...	4				
✖	S3 - Blocked Public Access (Bu...		Resource	4				
✔	S3 - Blocked Public Access (Ac...		Identity					
✖	[S3.4] S3 buckets should have ...	1/1	AWS Found...	4				
✖	S3 - Enabled Encryption At Rest		Resource	4				
✖	[S3.5] S3 buckets should requir...	1/1	AWS Found...	2				2
✖	S3 - Disabled HTTP Requests		Resource	2				2

The Sysdig platform discovers and presents you with a full inventory of your cloud so you can assess risk and compliance for assets, including IaaS, PaaS, hosts, containers, vulnerabilities, identities, and more. You can search and filter assets based on both static (e.g., public exposure, permissions) and dynamic risk factors (e.g., in-use packages).

Active cloud risk

CSPM requirements have shifted as cloud adoption has accelerated. The industry is moving beyond periodic posture checks to continuous posture assessments to identify, prioritize, and mitigate active cloud risks. Static checks are important, but given the speed of attacks in the cloud, these periodic point-in-time assessments can leave visibility gaps of hours or more.

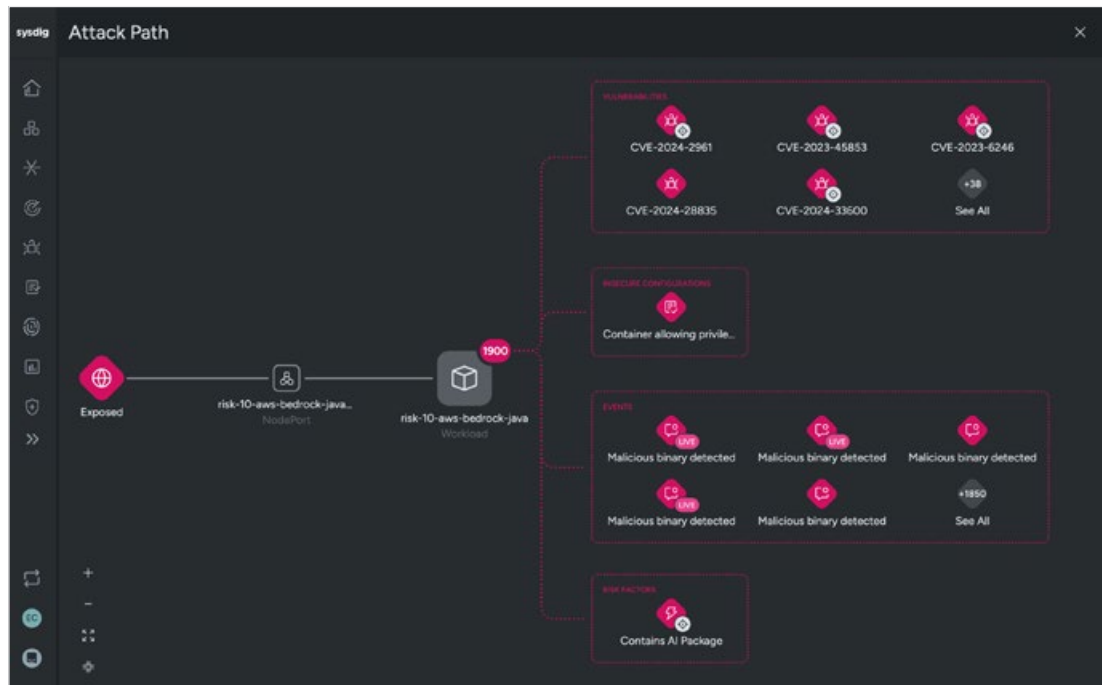
To spot active movement and dynamic changes, Sysdig enables the ability to spot and mitigate active cloud risk. Active cloud risk includes real-time activities and dynamic changes in your environment, such as:

- Risky identity behavior (e.g., user actively logging in with no MFA)
- Real-time configuration changes (e.g., connect to a known malicious network)
- In-use permissions (e.g., high-privilege access activated with no prior use)
- In-use packages with critical vulnerabilities (e.g., actively running software packages with high CVSS vulnerabilities)
- Workload threats (e.g., public encryption key uploaded)

Sysdig enriches static risk findings and overlays active risk information for prioritization, investigation, and remediation. The riskiest combinations of static and active risks to your AWS cloud are surfaced to the top with attack path visualization to speed up investigations. Guided remediation is integrated into the workflow to help security teams fix issues quickly.

By providing an [Attack Path Analysis](#) visualization, you can view interconnected risks and exploitable links across resources with active risks and events overlaid on static risks.

Figure 3:
Visualizing a cloud
attack path with
real-time insights



Infrastructure-as-code security

Sysdig maps assets and resources to your IaC manifest files to provide security insights, enable drift detection, and support the remediation of violations in your environment. For discovered violations, Sysdig generates tailored remediation suggestions that you can use to fix issues via pull requests through integration with your Git tooling.

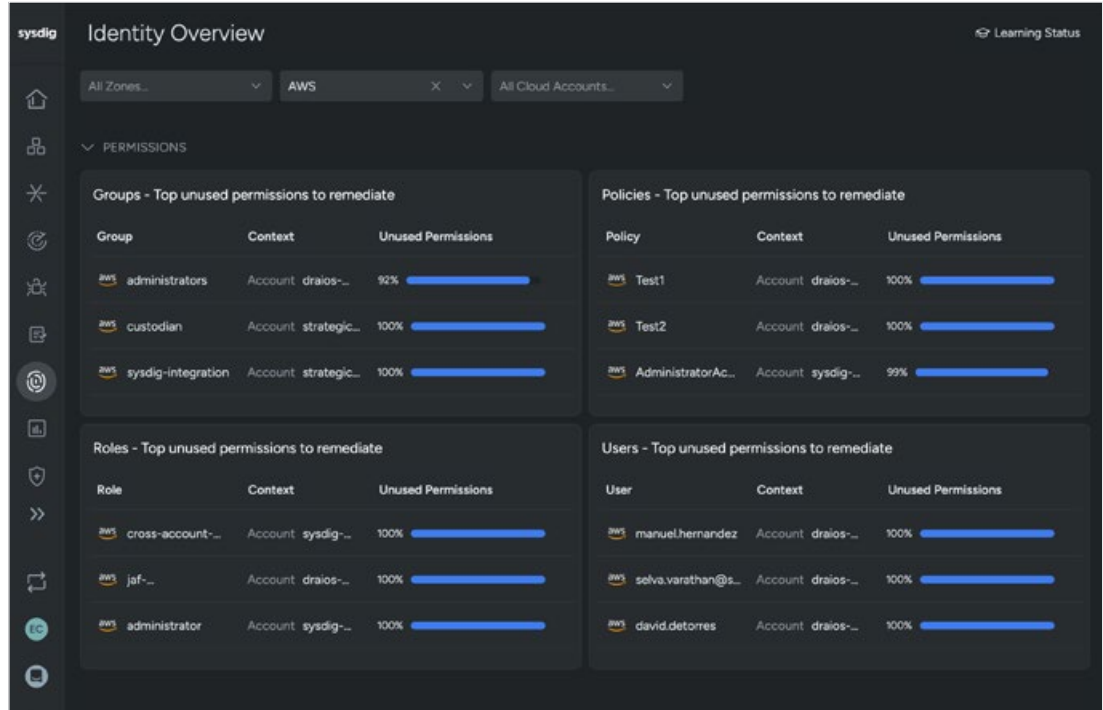
Entitlement management

Sysdig analyzes your cloud permissions to create a profile of your cloud users, roles, and policies. Audit logs analysis reveals the executed cloud commands in your AWS accounts and correlates this activity with policies, roles, and users. This visibility gives you insight into overly permissioned identities that may present the risk of credential misuse.

An identity and access dashboard informs you about:

- The total permissions given and used.
- How many users are inactive, and which users to consider deleting.
- The averages of permissions per policy and policies per user.
- The policies, users, and roles with the worst cases of unused permissions.

Figure 4:
AWS entitlement
and permission
management



Sysdig translates its permission analysis into policy suggestions that you can use to reduce excessive permissions and limit granted access to only what is needed.

To gain more insight into effective posture management for AWS, read [5 Essential Capabilities for a Modern CSPM Solution](#).

Cloud detection and response

Stopping cloud attacks is a critical capability needed for security teams as organizations continue to shift into larger and more complex cloud estates. Detection and response have been disrupted by noise and visibility gaps, often due to legacy EDR tooling. Cloud services, ephemeral containers, and identity sprawl create a dynamic and complex environment that can prove difficult to protect.

Threats can exploit weaknesses with minutes of exposure. Cloud Detection and Response (CDR) provides proactive defense against cyber attacks that target cloud infrastructure and data. It involves the continuous monitoring of cloud systems for potential threats, the assessment of severity, investigation capabilities, and the implementation of countermeasures to prevent or mitigate impact.

The key to effective CDR is the ability to identify, in real time, any malicious activity across workloads, identities, cloud services, and third-party apps to detect threats across the cloud fabric.

Microservice architectures running on containers, and orchestrated by solutions like Kubernetes, make applications faster to develop and easier to scale. However, monitoring container activity is exponentially complex. Containers may be distributed across multiple instances and hosts, and run programs in an isolated context. Getting visibility into activity requires unique instrumentation. Security instrumentation should be able to collect data in real time, but also shouldn't require modification to your container images to gain visibility.

Beyond technology, organizations also need to establish processes and ensure staff have the skills required to act swiftly to keep cloud systems and data secure.

With a sound CDR strategy in place, AWS users can:

- Reduce the risk of breaches
- Meet compliance requirements
- Reduce time to detect and respond
- Reduce cost, increase productivity, and securely accelerate innovation

One of the key security frameworks available for cloud teams to leverage to guide detection and response security strategies is the MITRE ATT&CK framework. It offers detailed, actionable information about attacker behaviors and techniques to help security teams proactively secure their cloud assets in an evolving landscape.

For a detailed overview, read the [AWS Cloud Detection and Response Matrix for MITRE ATT&CK](#).

10 Minutes to Pain

Targeted cloud attacks occur on average within 10 minutes of credential discovery.

AWS threat detection, investigation, and response solutions

AWS solutions for detection, investigation, and response enable you to detect cloud threats and gain a more complete understanding of the security of your workloads, applications, and data.



Amazon GuardDuty helps protect your AWS accounts, workloads, and data from threats. It monitors for malicious activity and delivers security findings to enable prompt response to potential threats or security incidents.



Amazon Security Lake enables you to store enriched multi-platform cloud security events on AWS where you can use your preferred analytics tools to analyze your security data.

Cloud detection and response for AWS with Sysdig

AWS cloud detection and response (CDR) capabilities from Sysdig focus on empowering analysts to guard against accelerated and complex cloud threats. Security teams gain deep visibility, context, and real-time detection capabilities built for the cloud and cloud-native workloads.

Built on [Falco OSS](#), Sysdig's CDR provides advanced detection and response capabilities across cloud logs, containers, Kubernetes, serverless computing, and cloud hosts. It detects threats in real time and correlates context across multiple domains to help analysts rapidly investigate, identify, and respond.

Security for AWS Cloud Services

AWS provides over 200 cloud services, including compute, storage, databases, analytics, networking, developer tools, management tools, security, and enterprise applications, to meet a wide range of use cases. To help you enable operational and risk auditing of your AWS accounts, AWS CloudTrail records actions taken by users, roles, and cloud services. It provides a key component for enabling security for AWS cloud services.

As your use of cloud services and infrastructure grows, Sysdig helps you automate the evaluation of CloudTrail events in real time using a flexible set of security rules. By continuously monitoring CloudTrail logs, you can detect and report suspicious cloud activity and events across a wide range of AWS cloud services.

Security for AWS Container Services

Sysdig has extensive expertise in container security and provides deep visibility to more easily detect threats, block attacks, and speed incident response for containers on AWS, including **Amazon Elastic Kubernetes Services (EKS)**, **Amazon Elastic Container Service (ECS)**, and **AWS Fargate**.

Sysdig employs kernel-level instrumentation to provide the most comprehensive runtime security for containers, including visibility into process creation, file system activities, network traffic, and more. It incurs a low overhead compared to user-space techniques, reducing impact on the overall system.

The Sysdig agent natively integrates into the container runtime API and Kubernetes API, enabling metadata collection and enrichment of generated events. Any event detected includes extensive information about affected assets, including the container name, Kubernetes cluster to which it belongs, pod, namespace, and service/deployment. Further enrichment is possible if you connect your AWS cloud account to correlate things like AWS account, cloud resource type, security group, and region.

By policy, containers can be stopped to immediately block a threat if specific activities are detected. In addition, a capture file that records all system activity from before, during, and after an event can be saved to enable post-event forensics and investigation.

Figure 5:
Kubernetes and
container security
overview



AWS Fargate Security

Serverless computing brings unique requirements for detection and response. Without access to a server host operating system, traditional agent-based instrumentation cannot be utilized. Sysdig uniquely solves the serverless visibility challenge for **AWS Fargate**.

- Serverless workload agents monitor each Fargate task for security events and enforce security and compliance policies.
- A serverless orchestrator agent collects information from serverless workload agents and sends it to the Sysdig SaaS backend cloud, so security teams can view events and take action.

Forwarding Security Events to SIEM and Security Data Lakes

Security information and event management (SIEM) and data lake solutions are used by security operations teams to store massive amounts of security-related data from various sources within an organization and are used for security monitoring, analysis, and compliance use cases.

Sysdig integrates with numerous SIEM and data lake solutions, including [Amazon Security Lake](#), enabling you to store enriched multi-platform cloud security events on AWS where you can use your preferred analytics tools to analyze your security data.

The 5/5/5 Benchmark for Cloud Detection and Response

Given the realities of the speed of attacks in the cloud, AWS users must measure their effectiveness in new ways. The 5/5/5 benchmark — five seconds to detect, five minutes to triage, five minutes to respond — challenges AWS users to acknowledge the realities of modern attacks and to push their cloud security programs forward. Achieving 5/5/5 requires the ability to detect and respond to cloud attacks faster than adversaries can complete them.

- **Detect threats within five seconds.** Organizations should be able to gather detection signals from their cloud security tools in real time to ensure visibility into ephemeral assets.
- **Correlate and triage within five minutes.** Teams should be able to gather full context for all correlated signals within five minutes of receiving the first relevant alert.
- **Initiate a response within five minutes.** Organizations should be able to initiate a tactical response within five minutes of confirming that an attack is in progress.

Visit the [5/5/5 benchmark page](#) to learn more.

Cloud security and generative AI

Generative AI (GenAI) is a top priority for organizations seeking to increase productivity and solve business problems. AI has the potential to aid cloud security by helping teams get a better understanding of risks and security issues, and even make security operations and response times faster. At the same time, organizations must find ways to manage the unique cybersecurity risks associated with operating GenAI and Large Language Model (LLM) applications.

AI security risks

GenAI presents great potential, but also comes with numerous security risks spanning privacy, cyber attacks, regulatory compliance, and breach of intellectual property. Some consider that AI may lower the barriers for threat actors to carry out sophisticated attacks and manipulate AI systems to compromise the system's integrity.

Because vast amounts of data are used and produced by AI, companies need to guard against issues such as unauthorized access and misuse, potentially breaching privacy regulations. To address these risks, enterprises must move toward maintaining the security, confidentiality, and integrity of AI, and carefully determine how best to prevent, detect, and respond to unauthorized access and adverse events.

Compliance frameworks for GenAI and LLMs

With AI becoming integral to countless aspects of business and society worldwide, concern over its impact on areas such as privacy, consumer rights, and national security has emerged. Governance frameworks and best practices intended to ensure the safe, private, and ethical use of AI are being developed to manage AI's expansive influence and mitigate associated risks.

Frameworks from organizations, such as [NIST](#), [Mitre](#), and [OWASP](#), seek to help those adopting AI guard against known risks and misuse. In addition, AI security regulations from around the world are in development or have already been implemented. These initiatives support a broader global trend towards addressing both the opportunities and challenges posed by AI.

Read [The Race for Artificial Intelligence Governance](#) to learn more.

AI as a security assistant

AWS users are looking toward Generative AI (GenAI) and Large Language Models (LLMs) to enhance security operations to prioritize risks, speed response, and simplify cloud security. AI can generate actionable insights that help teams get a better understanding of risks and security issues, and make security operations and response times faster. In addition, they hold the promise of enabling less experienced security personnel to handle complex tasks and improve overall cyber defenses through proactive risk investigation.

Sysdig security solutions and AI

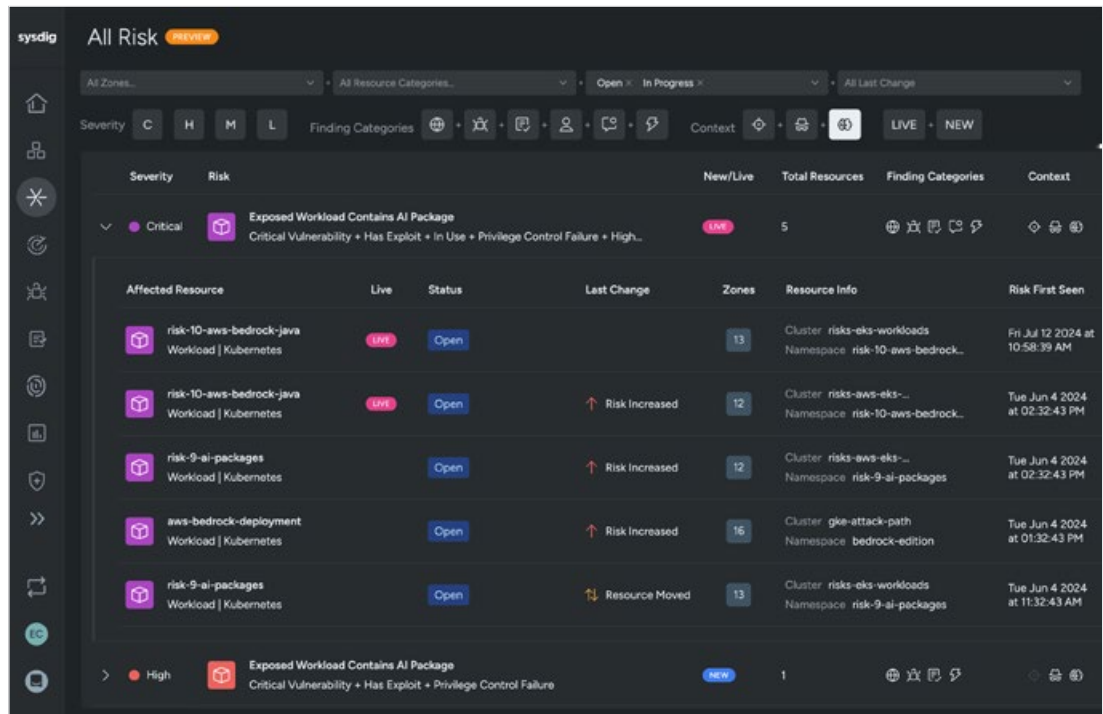
AI Workload Security with Sysdig

Sysdig provides AI Workload Security that helps companies securely adopt GenAI. It allows security teams to identify and prioritize AI workloads in their environment. This includes support for the leading AI engines and supporting software packages, including **Amazon Bedrock**, **Amazon Q**, and **Amazon SageMaker**.

AI Workload Security provides the visibility needed to establish data security measures that combat the risk of exposing trade secrets, proprietary information, and customer data through unauthorized access to AI workloads. A comprehensive view of correlated risks and events helps you quickly understand risk factors, including:

- Publicly exposed AI workloads
- In-use AI packages with critical vulnerabilities
- High-confidence threat events

Figure 5:
AI Workload
Security for AWS



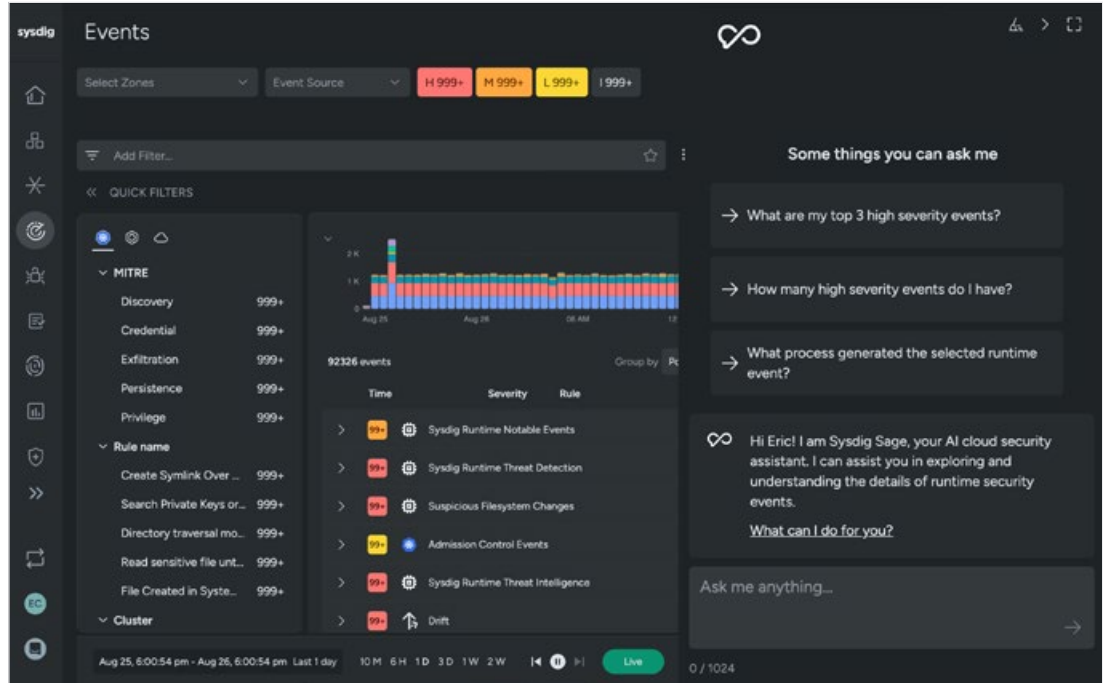
Learn more about [Sysdig's AI Workload Security](#).

Sysdig Sage™: The first AI cloud security analyst

Navigating the intricacies of securing public and private clouds, containers, and Kubernetes can be complex. Even seasoned professionals can find it challenging to stay ahead of the latest cloud threats.

Sysdig Sage is Sysdig's AI cloud security analyst. Sysdig Sage interacts with users through human-like conversations, helping to peel back the layers of security events. Sysdig Sage instantly delivers the collective wisdom of human experts and the continuous learnings of AI models to help AWS users accelerate their response to security issues.

Figure 6:
Sysdig Sage AI Cloud
Security Analyst



Learn more about [Sysdig Sage](#).

Conclusion

AWS is helping organizations move fast and innovate to deliver solutions that meet customer and market needs. As you scale your use of the cloud, your security practices must also evolve to incorporate measures to adapt to a growing array of threats. Robust cloud security enables real-time visibility and consolidates capabilities to correlate insights and accelerate response. CNAPP solutions, like the Sysdig platform, help AWS users modernize security and preemptively diffuse threats before they escalate and stay secure in the cloud.



In the cloud, every second counts. Sysdig stops cloud attacks in real time by instantly detecting changes in risk with runtime insights and open source Falco. We correlate signals across workloads, identities, and services to uncover hidden attack paths and prioritize the risks that matter most.

Sysdig. Secure Every Second.

[LEARN MORE](#)



sysdig

GUIDE

COPYRIGHT © 2020-2024 SYSDIG, INC.
ALL RIGHTS RESERVED.
GUIDE-010 REV H 9/24
