



sysdig

Securing AI: Navigating a New Frontier of Security Risk



HUMAN POWERED CONTENT

Alex Lawrence, Field CISO

Organizations worldwide are turning to artificial intelligence (AI) to gain insights, optimize operations, and remain competitive in an increasingly digitized economy. As businesses move workloads to cloud-native environments to support these AI initiatives, they encounter a new frontier of security risk. For security managers building cloud security programs, it's crucial to take a step back and ask: Does your current program truly cover the unique risks AI workloads bring?

Table of Contents

03	Introduction
04	The changing security paradigm — from castles to carnivals
06	Assumption of a breach
08	Building a security program for AI
09	Assess your security program for AI
11	Conclusion



Introduction

Organizations worldwide are turning to artificial intelligence (AI) to gain insights, optimize operations, and remain competitive in an increasingly digitized economy. By 2026, it's projected that more than 80% of enterprises will have adopted generative AI, integrating AI-enabled applications and APIs into their operations¹. Yet as businesses move workloads to cloud-native environments to support these AI initiatives, they encounter a new frontier of security risk. The dynamic, distributed, and ephemeral nature of cloud-native infrastructure combined with the complexity of AI workloads dramatically increases the attack surface.

AI workloads often interact with sensitive data, rely on APIs, and span multicloud and hybrid infrastructures. Each of these characteristics introduces potential vulnerabilities that adversaries can exploit. The result? Securing AI workloads calls for a fresh approach to security programs — one that reflects the complexities of cloud-native environments and the unique challenges AI introduces. While the principles of end-to-end security remain vital, they must now be complemented by frameworks tailored for these cloud and AI security use cases, such as those from MITRE ATT&CK and the Open Web Application Security Project (OWASP).

For security managers building cloud security programs, it's crucial to take a step back and ask: Does your current program truly cover the unique risks AI workloads bring? This means evaluating your practices against proven standards such as vulnerability management, runtime protection, and posture management. Beyond that, it's about equipping your team to perform deeper technical assessments to identify and address potential gaps that may expose AI workloads to risk.

This white paper explores how AI is changing the cloud security conversation and provides an assessment to help you evaluate your current security program to understand where it may fall short in addressing the unique needs of AI workloads. With these insights, you will be empowered to advance your security program, reduce risk, and enable your organization to securely leverage AI for innovation and growth.



By 2026, it's projected that **more than 80%** of enterprises will have adopted generative AI.

1 [Forbes, Adopting AI: From Interest To Action](#)

The changing security paradigm — from castles to carnivals

First, it's important to understand how the cloud has changed the security game. In the early days of enterprise IT, organizations built their security strategies around the “castle” model. A central data center functioned as a fortified stronghold with one clearly defined entry and exit point, guarded by firewalls, access controls, and intrusion prevention systems (IPSs). The focus was on protecting this perimeter, assuming that everything inside the castle walls was trusted and relatively safe. Security efforts focused on defending the network perimeter and monitoring inbound and outbound traffic to prevent unauthorized access.

However, the migration to cloud-native environments has dismantled the castle walls. Today's infrastructure resembles a bustling and dynamic carnival. The complexity of modern workloads — such as containerized applications, microservices architectures, and distributed environments — has dramatically transformed the security scope.

Workloads in cloud-native environments are often distributed across multiple data centers, hybrid environments, or even multicloud architectures. The dynamic nature of these environments means that workloads can spin up or down in seconds, making it almost impossible to predict exactly where a particular workload resides at any given moment. In fact, according to Sysdig research, 70% of containers live for less than five minutes,¹ which creates a high level of churn.

This shift dismantles the traditional perimeter, replacing it with a distributed network of constantly evolving, interconnected systems. APIs play a crucial role in this environment, enabling communication between microservices and third-party systems. Each of these APIs is also a potential entry point for attackers, expanding the attack surface and introducing new risks. Sensitive data flows across hybrid and multicloud ecosystems, spanning across services including virtual private clouds (VPCs), Simple Storage Service (S3) buckets, Relational Database Service (RDS) instances, even third-party systems such as identity management services. Securing this vast, interconnected landscape requires a shift to a more nuanced, cloud-tailored approach that focuses on securing data, services, and workloads wherever they reside.

AI workloads amplify security risks

The introduction of AI workloads further amplifies the security challenges in cloud-native environments. AI systems operate within a highly interconnected and dynamic infrastructure that brings unique risks. These systems often require vast datasets, many of which contain sensitive or regulated information, such as personally identifiable information (PII), financial data, or proprietary business intelligence. As AI becomes more deeply integrated into business operations, it also becomes an increasingly high-value target for adversaries.



20% of businesses have suffered an attack on their AI models in the past 12 months.³

The complexity of AI models themselves further complicates security. Many AI models are considered “black box” systems, where internal decision-making processes are difficult to understand and audit. This opacity creates significant challenges for security teams in detecting anomalous or malicious behavior, which puts organizations at greater risk of operational disruptions and security breaches that may go undetected until it's too late.

In Kubernetes environments, where workloads are containerized and distributed across clusters, this complexity is multiplied. Each AI workload running within a container may have its own communication channels, privileges, and dependencies — all of which must be secured. While providing orchestration for these workloads, Kubernetes also introduces unique challenges in managing access controls, network policies, and workload isolation. A breach within one container could potentially lead to lateral movement, allowing attackers to compromise additional workloads within the same cluster.

AI systems also rely heavily on APIs, external data sources, and shared infrastructures, further expanding the attack surface. For instance, an AI model that consumes data from external sources or interacts with other cloud services via APIs is only as secure as the weakest link in that chain. If an API or third-party service is compromised, attackers can potentially gain access to the entire AI ecosystem, including sensitive data and intellectual property.

This interconnectedness means that protecting AI workloads requires more than just securing the models themselves — it involves securing the entire ecosystem of services, data flows, and communication channels.

1 [Forbes, AI Models Under Attack: Protecting Your Business From AI Cyberthreats](#)

Assumption of a breach

In the cloud, businesses are navigating an ecosystem that is fast-paced, distributed, and inherently interconnected. With the modern cloud operating like a carnival, things are constantly changing: workloads spin up and down in seconds, APIs multiply, and sensitive data flows freely across hybrid and multicloud environments. When adding AI workloads to this equation — bringing their own reliance on sensitive data, APIs, and external integrations — the complexity and risks multiply.

Even organizations that have updated their security practices for the cloud must reevaluate their strategies to address the unique challenges posed by AI. At the core of this shift is adopting a security strategy that operates under the assumption of a breach. Why? Because in the cloud, malicious access isn't a distant possibility — it's an inevitability.

Why assuming a breach should be the default

In today's accelerated cloud environments, everything moves at rapid speed. A misconfigured API can be exploited within seconds, allowing attackers to move laterally or escalate privileges across systems in minutes. This pace makes the traditional "keep the attackers out" approach obsolete.

Cloud security isn't just about stopping threats before they happen; it's also about understanding that attacks that get past prevention controls are inevitable. By working from this understanding, organizations can ensure that they're ready to detect, contain, and mitigate threats before they escalate into significant incidents.

Of course, preventive controls are foundational — they're your helmet when you're skiing down the cybersecurity slopes. You wouldn't hit the slopes without it, just like you wouldn't deploy an application without making sure that basic security hygiene is in place. But just like a helmet won't stop you from falling, preventative controls won't stop every threat. You have to assume that at some point, an attack will get past prevention controls. And that's where detection capabilities step in to ensure that you can detect an attack the moment it happens.

Preventive controls provide a sense of safety, but they're inherently backward-looking. They focus on stopping what could happen, not what is happening right now. And with cloud attacks happening in seconds, the only way to stay ahead and know where to focus security investigations is with comprehensive cloud security that includes real-time threat detection capabilities.



Comprehensive cloud security

To effectively secure cloud environments, businesses need comprehensive cloud security capabilities that cover everything from data and applications to infrastructure and user access. This approach goes beyond simple defenses — it ensures that the entire cloud ecosystem is protected against evolving cyberthreats. By combining strong controls with continuous monitoring, comprehensive cloud security upholds data confidentiality, integrity, and availability, minimizing the risk of breaches and unauthorized access.

This strategy should span the entire cloud stack, from posture management to real-time runtime protection. This is where cloud-native application protection platform (CNAPP) solutions play an important role by covering the span of prevention to defense capabilities, such as:

- **Posture management**
Posture management focuses on continuously evaluating and enforcing secure configurations across your cloud environments. With real-time visibility into critical misconfigurations, you can take proactive steps to security risks, ensure ongoing improvements, and maintain a strong defense against evolving threats.
- **Runtime protection**
Cloud workloads need to be monitored in real-time to detect and mitigate threats as they occur. Through behavioral analysis and anomaly detection, runtime protection can actively contain threats during live operations, ensuring that malicious activities are stopped in their tracks.
- **Vulnerability management**
Cloud vulnerability management helps uncover blind spots by identifying vulnerabilities across your cloud environments and workloads — whether it's virtual machines (VMs), serverless functions, containers, or appliances. With continuous scanning, you can find, prioritize, and address critical vulnerabilities before they can be exploited.

By adopting comprehensive security across the entire cloud ecosystem, businesses can better defend against sophisticated attacks, adapt to evolving risks, and ensure that AI-driven innovations remain secure.



To effectively secure cloud environments, businesses need **comprehensive** cloud security capabilities.

Building a security program for AI

AI has quickly become a cornerstone of modern business — ignoring its role in your security strategy risks leaving security gaps unaddressed. As security leaders, the question you need to ask is: How can you support the organization's goal to leverage AI while ensuring an acceptable level of risk?

To ensure that your security program is aligned with modern needs, start by evaluating your AI usage. This requires asking the following questions:

- **Where are models being used?**
To effectively protect AI models, you must first know where they exist. This requires building a comprehensive AI inventory. You can't protect what you can't see, so identifying every instance where an AI model is deployed is critical to understanding its potential exposure and risks. Whether models are used internally or externally significantly influences their security posture and how best to monitor them.
- **Are they internal or external?**
Understanding whether AI models are deployed within internal systems or exposed to external environments is essential. Publicly accessible models increase the attack surface, as malicious actors can target them. In contrast, internal-only models may require a different level of access control and monitoring to ensure that only authorized users have the necessary permissions to interact with them.
- **What data do these models access, and how is that data protected?**
AI models rely on data, and knowing what data they access is fundamental to protecting sensitive information. You'll need to ensure that all data, especially proprietary or personally identifiable data, is encrypted and protected with the right access controls to avoid unauthorized exposure.
- **What permissions are granted to the models, and who is consuming them?**
Tracking who has permission to use AI models — and what level of access they have — is essential. With the cloud, it's easy to inadvertently grant broad permissions, so it's crucial to restrict access to only those who truly need it, and regularly audit these permissions to prevent privilege creep.

Answering these questions will provide clarity on your organization's current AI usage and security posture. Once you have a full understanding of where your AI models are operating, what data they access, and who has permission to use them, you can assess whether your current security program is sufficient or if it needs to be updated.

Building an AI security program isn't just about applying the basics of cybersecurity; it's about understanding the specific challenges AI introduces and adapting your strategy accordingly. If after evaluating your AI footprint you identify AI security gaps in your current cloud security strategy, this is a sign that your security program needs to evolve to keep pace with the demands of AI. Only by taking these critical steps can organizations safeguard their AI initiatives and minimize the risks associated with this powerful technology.

Assess your security program for AI

To determine whether your organization's security program is equipped to handle the unique risks posed by AI, start with this self-assessment. These questions are designed to help you evaluate your organization's current AI security readiness and identify areas that may require immediate attention.

For each question, answer **yes** or **no** based on your organization's practices.

01 Can you identify all AI models in your environment?

02 Are permissions for accessing your AI applications strictly limited to what's necessary?

03 Do you know which data sources your AI models can access?
Are they protected?

04 Are you actively monitoring for runtime anomalies?

05 Do you know if your AI workloads are publicly exposed?

06 Does your CSP have AI monitoring enabled by default?

07 Can you validate the security processes for pre-trained and third-party models prior to implementation?

08 Is access to underlying AI systems and training data restricted?

09 Are associated API keys securely managed and regularly rotated?

10 Do you have an incident response plan specifically for AI-related incidents?

Scoring

For each “no” answer, assign 10 points. For each “yes” answer, assign 0 points.

Now tally your points to determine your organization’s current level of AI security maturity. The more “yes” answers you have, the more aligned your security posture is with the demands of AI.

0-10 **Advanced program**
Your AI security program is mature.

11-20 **Moderate maturity**
Your program is on the right track, but you probably should update your security strategy to stay ahead of AI security risks.

30+ **High risk**
You should take immediate action to implement foundational AI security practices.

By completing this self-assessment, you’ll gain insights into where your organization stands and where improvements are necessary. If your results indicate a need for stronger AI security practices, our guide on adopting an AI security framework can help you get started on the path to protecting your AI-driven future.



Conclusion

AI workloads are driving innovation, but they also introduce growing risks. In cloud-native environments, where the pace of change is constant, securing these workloads requires a paradigm shift.

To ensure your security program aligns with modern needs, start by evaluating your AI usage. Ask critical questions about your AI footprint — where models are deployed, what data they access, and what permissions they have. This self-assessment will help you identify any gaps and determine if your current security strategy needs to evolve to address the demands of AI.

If your self-assessment reveals the need for an AI security program, it's time to take action. By combining strong preventive controls with real-time detective capabilities, organizations can build security programs that support strategic goals, while ensuring that AI drives business outcomes without introducing unacceptable levels of risk.

The stakes are high, but the rewards are greater. For security leaders, the question isn't whether to invest in AI security — it's whether they're ready to keep up with the pace of today's threats.

[CONTACT US →](#)

ABOUT THE AUTHOR



Alex Lawrence,
Field CISO at Sysdig

Alex has an extensive history working in the data center as well as with the world of DevOps. A majority of his career has been spent working in the world of OSS on identity, authentication, user management, and security.

