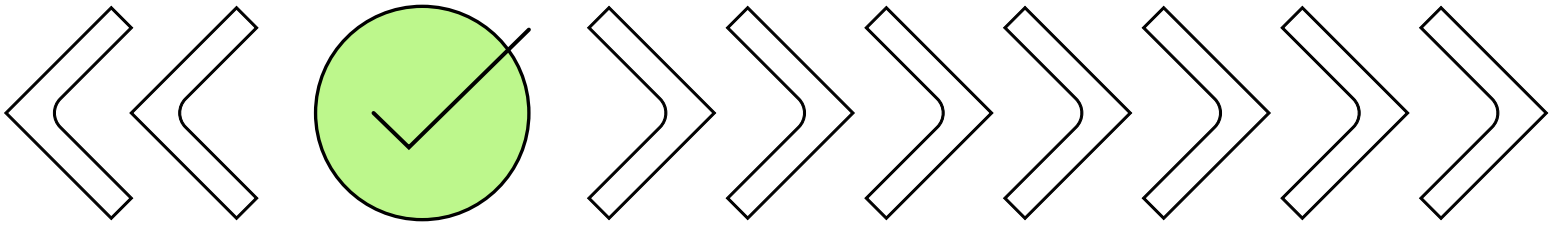
sysdig



CHECKLIST

Your checklist for securing AI workloads

Assess your AI security readiness in under five minutes

As AI adoption accelerates, cloud-native environments are dismantling traditional security perimeters. Dynamic, distributed workloads expand attack surfaces, while AI systems interacting with sensitive data across multicloud infrastructures introduce new vulnerabilities at every API and service boundary.

Permission sprawl compounds the challenge as containerized workloads multiply access points, and the ephemeral nature of cloud environments creates critical gaps in runtime visibility. Threats can emerge and disappear before traditional security tools can detect them.

This new paradigm is here to stay. Many organizations are encouraging or even requiring their employees to make greater use of AI, including asking engineers to use AI to write code, so as to take advantage of the speed and innovation automation enables. But that means teams must contend with these new security challenges.

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.	Yes	No
Can you identify all AI models in your environment?	◯	◯
Are permissions for accessing your AI applications strictly limited to what's necessary?	◯	◯
Do you know which data sources your AI models can access? Are they protected?	◯	◯
Are you actively monitoring for runtime anomalies?	◯	◯
Do you know if your AI workloads are publicly exposed?	◯	◯
Does your cloud service provider have AI monitoring enabled by default?	◯	◯
Can you validate the security processes for pre-trained and third-party models prior to implementation?	◯	◯
Is access to underlying AI systems and training data restricted?	◯	◯
Are associated API keys securely managed and regularly rotated?	◯	◯
Do you have an incident response plan specifically for AI-related incidents?	◯	◯
Do you have a complete audit log of who queried which model, with what inputs and outputs?	◯	◯
If using agentic AI, do you have visibility into what actions agents are permitted to execute, and are those boundaries enforced at runtime?	◯	◯

For each “no” answer, assign 1 point. 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.

Score:

0-2	**Advanced program** Your AI security program is mature.
2-3	**Moderate maturity** Your program is on the right track, but you probably should update your security strategy to stay ahead of AI security risks.
3+	**High risk** You should take immediate action to implement foundational AI security practices.

Sysdig helps security and development teams prevent, detect, and respond to threats instantly. No guesswork. No black boxes. **Just cloud security, the right way.**

REQUEST DEMO →

COPYRIGHT © 2026 SYSDIG,INC.
ALL RIGHTS RESERVED
CL-038 REV. A 09/26