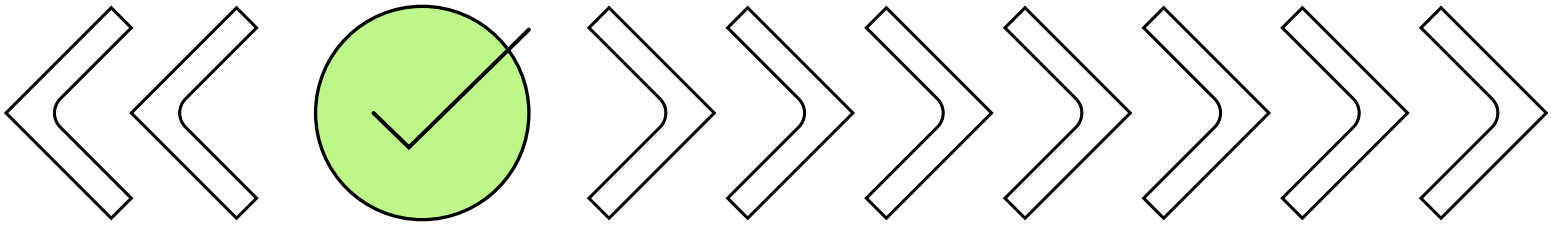
sysdig



ASSESSMENT

Your self-assessment for AI-powered cloud defense

Choose how you want AI agents to fit into your security workflows

Defending against machine-speed attackers takes machine-speed defense. But there is no single right way to bring AI into your cloud security program. Every organization sits at a different point on the AI maturity curve, with different tools, different workflows, and different risk tolerance. What matters is choosing the approach, or combination of approaches, that fits where your team is today and building toward where it wants to be.



Agentic AI

Agentic AI refers to AI agents built by a security vendor that reason over an environment, make decisions, and carry out multi-step workflows largely on their own. Rather than following a fixed script, an agent interprets what it finds, plans a response, and acts within boundaries a team sets ahead of time. Agentic AI fits teams that want the depth of a full security program without the overhead of managing the workflows themselves.



Headless cloud security

Some security teams already run their work through AI coding agents like Claude, Cursor, or Codex, and expect their tools to meet them there rather than pulling them back into a separate interface. Headless cloud security is built for exactly that scenario.

Instead of a UI, headless security exposes runtime-grounded data, context, and decision logic through APIs, a CLI, and MCP services, packaged as curated agent skills built from more than a decade of cloud security research. Headless security gives teams that are all in on AI a way to build security into the workflows they already use.



Chat-based assistants

The first wave of AI for security took the form of a chatbot layered onto an existing platform, built to answer questions in plain language and point a user toward the next step. For a team just starting to bring AI into its security program, that's often enough: Plain-language answers and navigation help without changing how the team already works.

The right starting point depends on where your team is today. Answer the questions below to see which approach, or combination of approaches, may fit best.

For each question, choose the answer that best describes your team:

01 Where does AI fit into your security work today?

- ◯ A. We're still evaluating what AI can do for our team.
- ◯ B. We use AI in some daily workflows and want to automate more.
- ◯ C. We work through AI coding agents (e.g., Claude Code, Cursor) in our everyday work.

02 When AI helps with a security finding, what do you want it to do?

- ◯ A. Answer our questions and recommend next steps, while our team does the work.
- ◯ B. Take on the triage and follow-through in a guided workflow, checking in on key decisions.
- ◯ C. Work inside our own tools and workflows, such as drafting a fix for approval as a pull request in our repo.

03 Where does your team want to investigate and fix security issues?

- ◯ A. In the security platform we already use, with AI answering questions along the way.
- ◯ B. In a guided workflow inside the security platform that walks us through each step.
- ◯ C. In our own AI coding agents and tools, alongside the rest of our engineering work.

04 How do you want to shape security workflows?

- ◯ A. Use what's already built, since we aren't ready to design on our own.
- ◯ B. Let AI agents in the platform run guided workflows for us.
- ◯ C. Customize our own workflows and connect them to our tools and data.

05 What best describes your team's experience with AI?

- ◯ A. Mostly chat and question-and-answer.
- ◯ B. Some automation in specific workflows.
- ◯ C. Agents we configure and operate ourselves.

06 What is your biggest challenge in security today?

- ◯ A. We spend too much time working out what findings mean and where to start.
- ◯ B. Findings pile up faster than we can triage, assign, and track them.
- ◯ C. Fixes stall because teams work in different tools and turning findings into action is manual.

Scoring

Count how many times you chose each letter. Your most common letter suggests where you may be today and your best starting point.

A: B: C:

Mostly
A

Exploring AI

You are probably still evaluating what AI can do for your team and may be looking for clearer answers and guidance. Basic AI assistance, such as a conversational AI assistant or chat experience, could be a good starting point. This would give you plain-language explanations of risk and suggested next steps, without major changes to how your team works today.

Mostly
B

Working with AI

You are probably using AI in some workflows and may be ready to automate more of the day-to-day. Agentic AI could be a good place to start, with AI agents running guided workflows in your security platform.

Mostly
C

Building with AI

You are probably already working with AI coding agents and may want security to show up there. Headless cloud security could be a strong fit, bringing security data and expertise into those tools so you can customize workflows around your own tools and data.

If your answers are split, or two letters tie, you may be somewhere between stages, and a combination could fit best. These approaches work together, so you can use each one where it fits the task.

COPYRIGHT © 2026
SYSDIG,INC.
ALL RIGHTS RESERVED
CL-042 REV. A 10/26

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 →