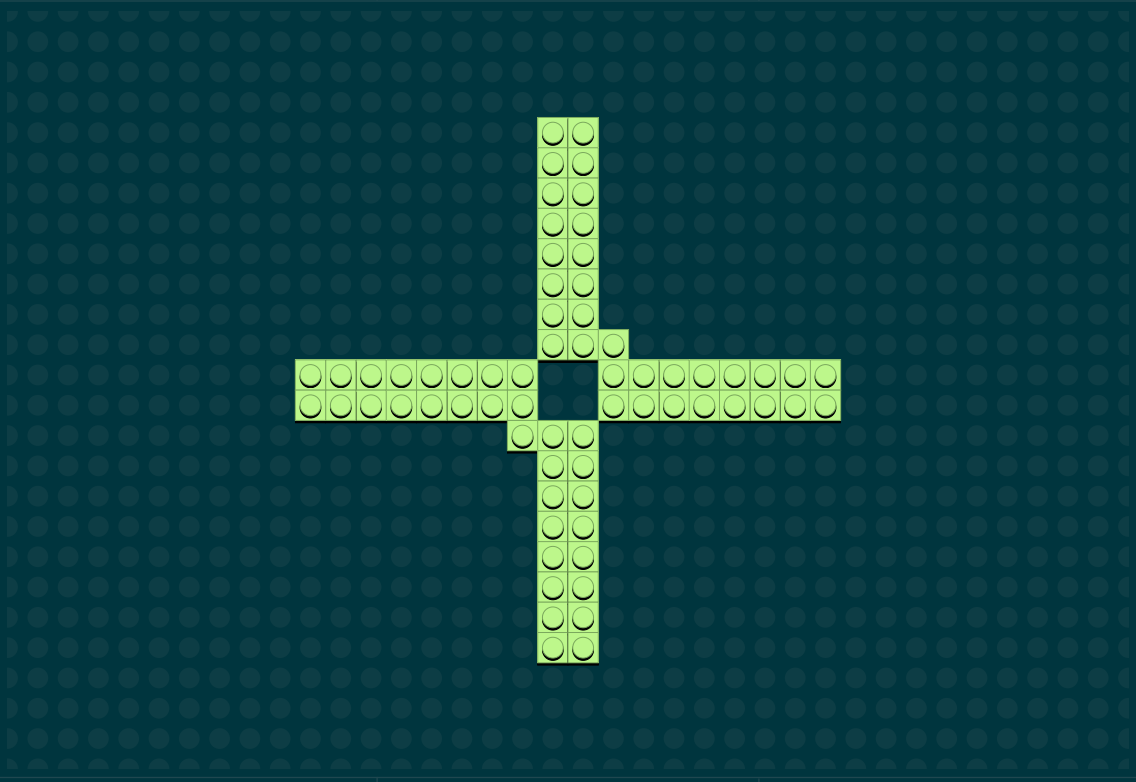
sysdig



Your Blueprint to

AI-Powered Cloud Defense, the Right Way

Attackers already have AI agents working at machine speed. This blueprint breaks down three ways to give your defense the same leverage (agentic AI, headless cloud security, and chat-based assistants), what to demand from a vendor, and a self-assessment to find your starting point.

Contents

03
The AI advantage

05
The problem: Defense built for a slower world

07
Putting AI to work in your security program

12
What to look for in an AI vendor for cloud security

15
The results: From faster defense to better outcomes

17
Self-assessment: Where does your team stand?

20
Conclusion: Defense that moves at machine speed

The AI advantage

Attackers have gained a decisive new advantage in the cloud: machine speed. Powered by AI agents, modern threats can discover vulnerabilities, chain familiar techniques, and adapt on the fly without waiting for human instruction. Defenders, meanwhile, still largely operate in human-speed cycles, resolving incidents one investigation at a time.

At Sysdig, we believe AI for cloud security should give defenders the same kind of leverage. Your team stays in control, setting the goals and making the decisions that matter, while AI takes on the work that slows you down: finding what matters, understanding how risks connect across the environment, and acting quickly when something needs attention.

Done right, AI-powered cloud defense can help your security team move from constantly catching up to keeping pace with an evolving threat landscape. It can reduce the manual work between discovering a threat and doing something about it, while helping you protect increasingly complex cloud environments without adding more operational burden.

This blueprint lays out what that looks like in practice: how AI agents can help build a more responsive approach to cloud defense.

7 days



Attackers are using AI to move faster, scale attacks, and exploit vulnerabilities at machine speed. Exploits now occur, on average, seven days before a vendor patch is available.[1]

1 Mandiant M-Trends 2026 Report

The problem: Defense built for a slower world

In early 2026, the AI advantage was proved on three separate fronts within months of each other: the vulnerability discovery process, a ransomware operation, and a major AI infrastructure provider. Each time, AI moved too quickly for human-centered workflows to keep up.

The vulnerability discovery front

The first front emerged in April, upstream of any attack: finding the vulnerabilities in the first place. Anthropic's Claude Mythos Preview generated 181 working exploits against Firefox's JavaScript engine, compared to two from the prior frontier model, surfacing a 27-year-old bug in OpenBSD and a 16-year-old flaw in FFmpeg that five million rounds of automated fuzzing had missed.[2] The significance is less about any single exploit than the shift in how quickly frontier models can identify and validate vulnerabilities across large and complex codebases.

The ransomware front

The next came in July, when the Sysdig Threat Research Team documented the agentic threat actor (ATA) JADEPUFFER, which is the first documented fully agentic ransomware attack on record. An LLM agent exploited a known Langflow vulnerability for initial access, harvested credentials, moved laterally to a production database, and executed a self-narrating extortion operation across more than 600 actions, encrypting more than 1,300 configuration records.[3]The technique wasn't new, but the speed and consistency of an agent running it end to end was.

The infrastructure front

That same month, on a different front, Hugging Face disclosed that an autonomous AI agent had breached its production infrastructure. The agent exploited two code-execution flaws in Hugging Face's dataset pipeline, escalated privileges, harvested credentials, and executed tens of thousands of automated actions across short-lived sandboxes. Investigators later traced it to OpenAI models running an internal security evaluation, rather than a threat actor.[4]

2 **Sysdig, Anthropic Mythos just broke the four-minute mile in cyber offense, April 2026**
3 **Sysdig, JADEPUFFER: Agentic ransomware for automated database extortion, July 2026**
4 **Embrace The Red, Autonomous AI Intrusions Are Here: Lessons from the Hugging Face Compromise, July 2026**

Threat intelligence confirms this pattern is already widespread.

The significance of these examples extends beyond the individual incidents. They reflect a broader shift in the speed and scale of modern attacks, one that threat intelligence is beginning to quantify.

Active exploitation now begins an average of seven days before a vendor patch is available, meaning attackers move against a flaw before defenders have anything to patch.[5] A human analyst working through alerts one at a time, however capable, can't investigate, decide, and act as continuously as an agent that never stops. Matching that pace in the cloud requires AI-powered cloud defense.

5 **Mandiant M-Trends 2026 Report**

Pillar

1

Putting AI to work in your security program

Defending against machine-speed attackers takes machine-speed defense. The goal is a security program that moves at that pace, while your team sets the direction and AI agents handle the volume of work no human team could keep up with alone.

Done right, AI changes how that work gets divided. Your team approves decisions that carry weight, and stays in control of what happens in your cloud environment. Your agents handle the work that consumes time: finding what matters, gathering the context, taking action where they are authorized to do so, and learning from what they encounter. That shift gives you a way to close the gap between the speed of attack and the speed of defense, so your business can move quickly without taking on more risk.

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. Two approaches make up the core of that choice.



> The reality is simple: Security teams don't need more dashboards; they need better outcomes.

Loris Degioanni
CTO and founder, Sysdig

Agentic AI: A guided team of specialists

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.

At its best, agentic AI will design a full workflow to achieve the outcome you want, such as reducing critical vulnerability exposure or closing posture gaps in production, and lets specialized agents build the plan, gather the context, and carry it out. One agent might prioritize vulnerabilities against real runtime exposure, another investigate suspicious activity, and another manage posture and compliance drift. Together, they compress work that used to take a team days into a process measured in minutes.

Agentic AI fits teams that want the depth of a full security program without the overhead of managing the workflows themselves. Not every agentic approach gives teams the same control. Ideally, agentic AI should require human approval before high-impact changes reach production and record every action it takes for review. With these controls in place, it works well for organizations still building confidence in AI-driven workflows.

How this plays out

An exposed Langflow instance goes live in a development environment.

Before an attacker finds it, an AI agent flags the critical exposure, checks the CVE against the environment, and confirms it's exploitable and internet-accessible. The agent then opens a ticket with a one-click fix and routes it to the team that owns it.

Before the fix is approved, suspicious database activity surfaces.

Instead of a multi-hour manual dig through logs, an investigation agent traces the exposed instance to the database it touches in minutes, following the same path an attacker would have taken. The team approves the fix and closes that path before it turns into the next JADEPUFFER.

Headless cloud security: Bring your own agent

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. Rather than teaching a general-purpose coding agent to reason about cloud risk from scratch, headless security hands it the expertise directly: how to triage a vulnerability against runtime exposure, spot a misconfiguration and propose a fix, and correlate a runtime alert with the vulnerabilities and identities around it.

Your team can then build security into the workflows it already uses. Because headless security isn't tied to vendor-defined workflows, teams can shape security around the way their business operates. Fixes can be opened as pull requests alongside the code they affect. Reports can be generated on demand with a plain-language prompt, and alerts can route wherever your team already works, instead of sitting in a console nobody opens until something breaks. Headless security gives teams that are all in on AI a way to make security part of the workflow.

Headless cloud security in practice

A few examples of how this changes the workflow:

- Ask a coding agent to triage a vulnerability in a running cloud workload. It pulls real-time runtime context and opens a pull request.
- Catch a misconfiguration, such as an exposed storage bucket or an overly permissive IAM role, before it turns into a real exposure. The agent proposes the fix inline, where your infrastructure code already lives.
- Request a cloud posture report from wherever your team already works. Get one back in seconds, scoped to what's running in production.

A word on chat-based assistants

Before agentic AI and headless security, the first wave of AI for security took a simpler form: 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.

As teams grow more comfortable putting AI to work, most move from asking questions toward letting agents act directly, whether inside a guided platform experience or through their own coding agents.

Pillar

2

What to look for in an AI vendor for cloud security

AI doesn't make a cloud security program effective simply because it is present. The real question is whether the AI can see what's happening, apply real judgment to it, and act inside limits your team controls. Those foundations determine whether AI can move beyond generating answers and genuinely help defend a live cloud environment.

When evaluating a vendor for AI cloud defense, look for capabilities that support four things:

1. Deep runtime intelligence and telemetry

AI is only as good as the data it runs on. An agent reasoning over stale scans or point-in-time snapshots will make the wrong call with complete confidence. Look for a foundation built on real-time, kernel-level visibility: what is running right now, separate from what a configuration file says should be running. That distinction is what separates a genuine risk from a theoretical one.

2. Vendor cloud security expertise

A model is not a security analyst. The vendors delivering real value have spent years studying how cloud attacks unfold in practice, and they package that expertise directly into their agents rather than asking a general-purpose model to figure it out from scratch. That experience is what turns a fast AI agent into a skilled one.

3. Adaptive, personalized security

No two environments carry the same risk profile, and a one-size-fits-all workflow will always force some part of your team to work around the tool instead of with it. The right AI vendor lets your organization define the outcomes and priorities that matter to your business, then adapts its agents to fit, learning your environment more precisely the longer it runs.

4. Built-in accountability and guardrails

Speed without oversight is its own risk. A trustworthy AI vendor keeps a human in the loop for the decisions that matter, requiring approval before high-impact actions land in production. It records what every agent did and why, tracks progress against real metrics, and gives your team a clear answer when the board asks whether security is truly improving, beyond a count of alerts reviewed.



The payoff shows up in scale. With Sysdig Secure AI, Sysdig's own AI-powered cloud defense solution, teams can handle more than 10x as many investigations as human experts alone could manage, at 88% lower cost.

10x

more investigations

88%

lower cost

Pillar

3

The results: From faster defense to better outcomes

Putting AI to work for your cloud security changes what your team can accomplish. You can respond faster, cover more of your environment, and spend more of your time on the decisions that require human judgment.

Reduce risk faster

AI can identify, investigate, prioritize, and help remediate security risks faster than human teams can alone. Exposure windows shrink, threats get contained sooner, and the gap between finding a risk and resolving it gets smaller.

Make security fit your business

AI can adapt security to your environment, priorities, and existing workflows rather than forcing your team into a predefined process. A good AI solution builds context about all three over time, helping make security increasingly specific to how your business operates.

Give your team time back

The repetitive work of investigation, ticketing, reporting, coordination, and more can move to AI, freeing your team to spend more time on the decisions that require human judgment. AI can take on that work while keeping analysts in control.

Cover more of what matters

AI can take on work your team might otherwise have to leave for later, so more investigations, exposures, and alerts get the attention they need across your environment. The result is broader coverage without requiring your team to scale at the same rate.

See your security improve

AI can help turn security activity into measurable outcomes. From coverage and exposure time to SLA adherence, AI can track the metrics that matter, so when leadership or the board asks whether security is improving, you have evidence to show them.

Self-assessment: Where does your team stand?

AI-powered cloud defense can be introduced in more than one way, and 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.

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.

Conclusion: Defense that moves at machine speed

AI-powered attacks have changed the pace of cloud security. Adversaries can discover vulnerabilities, connect attack paths, and execute familiar techniques at machine speed. Defenders need a cloud security model that can keep pace.

AI-powered cloud defense gives security teams that leverage. Agentic AI can investigate, plan, and execute security workflows inside a UI. Headless security can bring cloud security expertise into the AI workflows and tools teams already use. Chat-based assistants can help teams take their first steps toward AI-driven security. The common thread is the same: AI takes on more of the work while your team stays in control.

That requires more than adding an AI layer to an existing security tool. Effective AI-powered cloud security needs the right data, deep cloud security expertise, runtime context, and guardrails that make every action accountable. It needs to turn AI from something that answers questions into something that helps your security team achieve measurable outcomes.

That is the approach Sysdig delivers with Secure AI. Grounded in runtime truth and Sysdig's cloud security expertise, Secure AI gives agents the context to understand what is happening in your environment, determine what matters, and take action within the boundaries your team sets.

The result is security built for the way cloud environments operate today: AI works at machine speed. Your team stays in control. And defense keeps pace with the speed of attacks.

Sysdig helps security and development teams prevent, detect, and respond to threats instantly. Founded by the creators of Falco and Wireshark, and built on agentic AI, Sysdig delivers real-time defense grounded in the uncompromising truth of runtime. No guesswork. No black boxes. Just cloud security, the right way.

GET A DEMO

ABOUT THE AUTHOR



Matt Kim,
Senior Product Marketing Manager at Sysdig

Matt is a product marketer with a focus on cybersecurity, particularly in container security and vulnerability management. With over a decade of experience, he brings keen insight into translating complex security topics into straightforward, actionable insights that empower teams to tackle today's toughest security challenges.

sysdig

BLUEPRINT

COPYRIGHT © 2026 SYSDIG,INC.
ALL RIGHTS RESERVED.

BP-007 REV. A 10/26