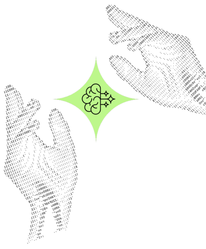
sysdig

Your visual guide to

Headless cloud security

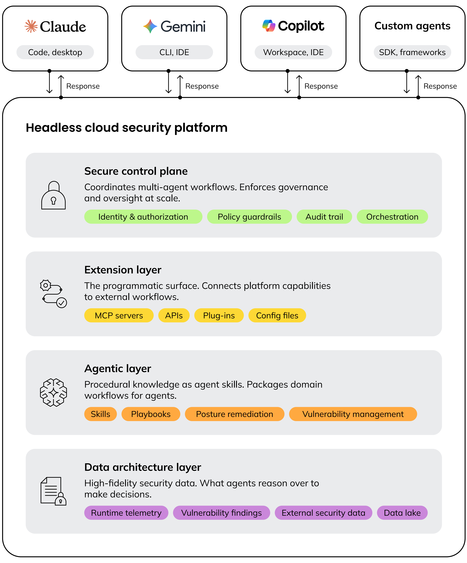


Most cloud security workflows revolve around a dashboard: A person logs in, reviews alerts, and decides what to do. But with attacks unfolding in minutes, that model can't keep up.

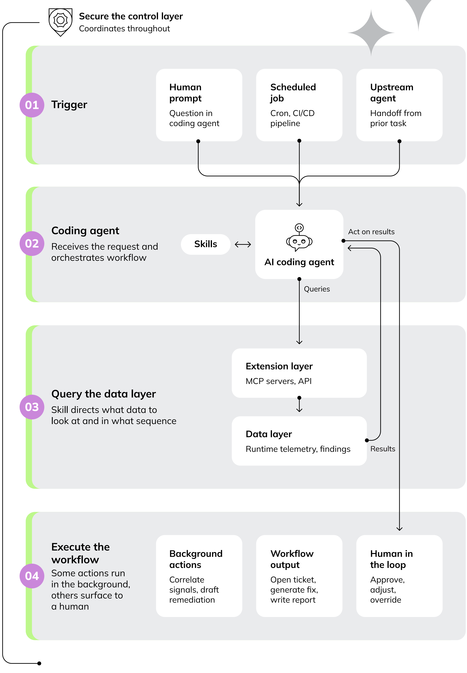
Headless cloud security drops the dashboard. Instead, AI agents run security workflows through APIs while humans set policy and approve key actions. This also meets engineers in the AI coding assistants where many of them already work.

In this infographic, we'll take a look at the architecture and workflow of headless cloud security, and what to consider when adopting this new approach.

Core components of a headless cloud security architecture



Headless cloud security workflow in practice



Tips for headless cloud security adoption

01

Define where agents can act

Decide which actions agents take alone, which need approval, which need human authorship.

02

Treat agents as managed identities

Every agent needs scoped permissions, audit trails, and lifecycle controls.

03

Protect the integrity of skills and inputs

Treat them as code: version-controlled, reviewed, validated.

04

Decide what gets persisted in memory

Agent memory is powerful when trustworthy, a liability when it isn't.

05

Build observability for agent activity

Tool invocations, data queries, decisions, actions. Log the pattern that actually exists.

06

Start narrow, then expand deliberately

Validate behavior in constrained scope before extending it.

sysdig

INFOGRAPHIC

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

Want to learn more?

Discover the Secure AI platform for AI-powered cloud defense, the right way.

LEARN MORE →