INTEGRATING THE CLAUDE COMPLIANCE API

Every asset, now including Claude.

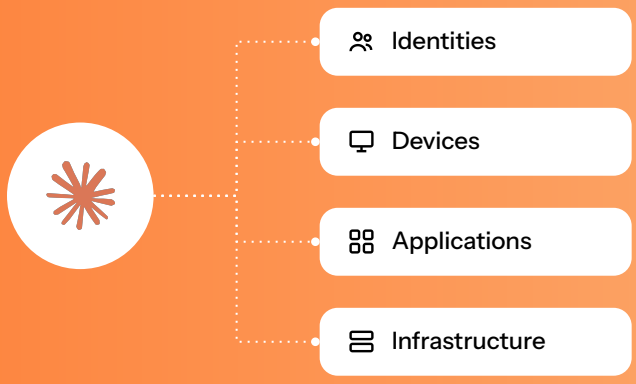

Axonius brings your Claude Enterprise footprint into the Asset Cloud, reconciled against the rest of your environment to surface the shadow AI, risky access, and coverage gaps that no single tool can see.

Organizations are adopting Claude across the enterprise, and its footprint touches your SaaS, software, and identity estate in new ways. **Safe AI adoption demands decision-grade asset intelligence as the foundation.**

Claude doesn't have to be a blind spot. Reconciled into the Asset Cloud, every user, account, and key resolves to the person behind it, their security posture, and the rest of the access they hold. That's decision-grade asset intelligence on Claude, held to the same standard as everything else you run.

HOW THE ANTHROPIC ADAPTER WORKS

01 Authenticate

The adapter connects to the Claude Compliance API with read access to your Claude Enterprise workspace.

02 Fetch assets

On a configurable schedule (default daily) it pulls users, groups, roles, permissions, and API keys.

03 Reconcile

Each artifact lands as a first-class asset, normalized and related to identities, devices, software, and cloud.

04 Assess

Axonius surfaces unmanaged Claude users, accounts on devices without EDR, excessive access, and more.

Net-net: your Claude footprint is fully modeled against everything else in your environment, ready to discover, prioritize, and act on with all the right context.

DISCOVER SHADOW AI

See the Claude usage your tenant can't report

The Compliance API tells you the managed users in your tenant. Axonius discovers Claude and Claude Code from the software installed across your devices, then compares it against that population. The negative space is your shadow AI — people running Claude on company machines who never appear in the tenant. Because reconciliation already ties each person to a device and identity, that gap becomes a ready-made migration list.

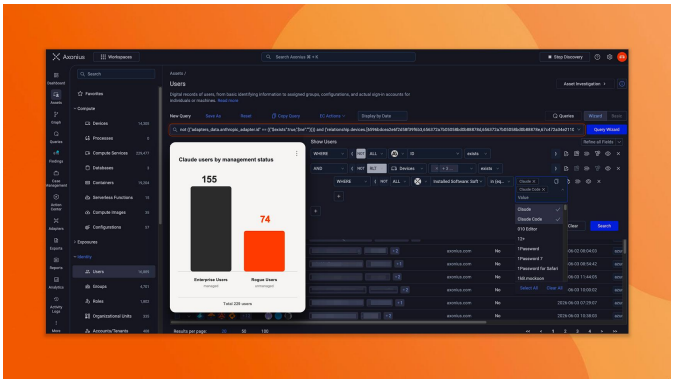

find unmanaged Claude users to migrate

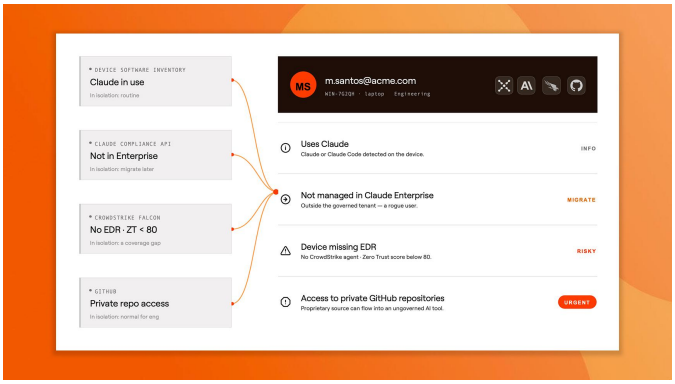

surface the toxic combinations that indicate risk

REDUCE THE ATTACK SURFACE

Run Claude through the same risk model as everything else

Findings that matter rarely come from a single fact. On one engineer's laptop, four unremarkable signals — Claude in use, an account outside the tenant, a device without EDR, access to private repositories — climb from informational to urgent once reconciled onto a single asset record. Each notable finding is tied to the owner who can fix it, with the full context to make a decision.

Safe AI adoption is a *fundamentals problem* before it's an agents problem.

Every tool you add widens the attack surface, and AI is widening it fast. That puts a premium on the fundamentals — a trustworthy inventory, closed coverage gaps, every asset seen in context. Your AI strategy starts from ground you can see and trust.

GETTING STARTED

Add the Anthropic adapter in your Axonius instance and configure read access to your Claude Enterprise workspace. Full setup lives in the Axonius documentation.

Request a demo →