




How KeyValue helped built Axari's AI Twin, a Claude-powered security teammate

Report type
Case study

Domain
Cybersecurity · AI agents

Headquarters
🇺🇸 San Jose, California



CLIENT REQUIREMENT

A teammate that does the security work, not another dashboard

Axari was founded by Prasen Shelar, the first product manager at Demisto (later acquired by Palo Alto Networks). His belief was simple: security teams don't need more alerts. They need someone who can act on them and see the work through.

Axari partnered with KeyValue to bring that vision to life: an AI Twin for every security professional.

THE CHALLENGES IT NEEDED TO SOLVE



More alerts than teams can handle

Security tools surface issues faster than teams can investigate, prioritise, and assign them.



Too much chasing, too little fixing

Much of a security lead's week goes on status updates, reminders and audit evidence.



Closed tickets, unresolved risks

Marking a ticket as done doesn't guarantee the underlying security issue has been fixed.

OUR ROLE

Building alongside Axari’s team

KeyValue worked with Axari’s team to build the platform and AI Twin, with Claude at the centre from the very first commit, and continues to ship new security workflows.

THE CORE PRODUCT

SUPPORTING THE TWIN

 Axari AI Twin The security teammate itself: reasoning, multi-step workflows, memory and proactive intelligence on Claude, working inside Slack, Teams and other work platforms.	 Configuration dashboard Where teams set up their Twin, connect integrations and manage settings	 Integrations 50+ security, identity and productivity tools	 Safety & trust Approvals, message checks and a full audit trail	 Cloud & operations Secure hosting, agent sandboxing and monitoring
---	---	--	---	--

WHAT WE BUILT

An AI Twin that carries security work from alert to verified fix

One Twin per security professional, not another tool. It lives in Slack, Teams and the other platforms teams already work in, learns the people, systems and priorities around its owner, and handles the legwork while they keep the decisions.



01 · INVESTIGATE

Finds what matters

Pulls findings from security scanners and alerts, assesses exploitability and asset ownership, and prioritizes the issues that need attention first.



02 · ACT

Does the work in your tools

Creates and assigns tickets, prepares fixes, and gathers audit evidence across 50+ tools, operating within the owner's permissions.



03 · FOLLOW UP

Follows through until it's done

Reminds owners, escalates stalled tasks, and increases urgency as deadlines approach, without sending duplicate reminders.




04 · VERIFY

Confirms the fix worked

Closes an issue only after a fresh scan verifies that the vulnerability has been resolved, then reports back to the owner.

WHAT WE BUILT

One pattern behind every workflow



For a security team to trust an AI agent, it needs more than the ability to act. It needs accountability. That principle shaped how we built the Twin: operating within the owner's permissions, requiring approval for high-risk actions, and keeping every action traceable.

FEATURES IN SPOTLIGHT

Built for how security teams actually work



60+ ready-made tasks & workflows

From a quick risk summary to full vendor onboarding.



Automations in plain language

Describe a recurring job in a chat message, and the Twin runs it.



Remembers people & ownership

Knows who owns what, and what was already discussed.



Daily morning brief

The few things that matter, from every connected tool.



Proactive intelligence

Surfaces slipping work, alerts and unprepared meetings before anyone asks.



Human in the loop

Risky actions need approval by default, and teams can configure which ones.

CHALLENGES WE SURPASSED

Making an AI agent safe to run on real security systems

Getting an AI model to respond is the easy part. The real engineering challenge was making the Twin safe and reliable enough to act within a customer's systems.



Acts only within authorized limits

The Twin uses the owner's existing permissions and seeks human approval before taking high-risk actions, keeping people in control.



Keeps sensitive information protected

Every outgoing message is checked for secrets and personal data, with sensitive information removed or corrected before it leaves the system.



Long-running work that keeps its place

Security workflows can run for days across many steps and people. If a system restarts midway, the Twin resumes from the last completed step instead of starting over.



One customer. One secure environment

Each customer's Twin operates in its own isolated environment with dedicated credentials, keeping access and operations separate.

TECH STACK

Built on a modern, Claude-first stack

AI

 Claude · Opus, Sonnet, Haiku

 Claude Agent SDK

 Model Context Protocol

 Honcho memory

 Qdrant

PRODUCT

 React

 Python

CLOUD

 AWS

 Kubernetes

 Agent sandboxing · Tensorlake

QUALITY

 LLM tracing · Netra

 OpenTelemetry

 VictoriaMetrics

 Grafana

 Audit trail

BUILT WITH CLAUDE · ANTHROPIC

The right Claude model for every job



Claude is the default model for the Twin’s core work

IT DRIVES THE TWIN’S

- Reasoning
- Understanding of requests
- Choice of tools
- Drafting

Claude Opus, Sonnet & Haiku

Matched to each job, from multi-step security workflows to fast checks on every outgoing message.

Agent SDK & MCP

Secure, permission-checked access to connected tools.

Prompt caching

Implemented custom prompt caching strategies on the Claude Agent SDK, cutting the cache-write premium by 75% without compromising speed.

RESULTS

What an AI twin changes

10 days
→ 30 min

From security finding to assigned ticket

2 weeks → 30 min

Quarterly audit preparation

4 hrs
→ 20 min

Daily first-line alert
triage

Next audit
→ same day

Spotting drift in
security controls