

SOLUTION BRIEF

Axonius has the data *Axonius AI* has the answer

Your environment has never been more complex, and the answers you need have never been more scattered.

- Which high-risk assets are missing critical updates?
- Where is shadow AI running?
- Is that spike in device count real, or a data quality problem?

Axonius AI turns a plain-language question into an answer grounded in your cyber asset data. It reasons across devices, users, software, SaaS, identities, cloud resources, and exposures, all correlated into one reconciled model. And with your approval, it acts. **The constraint was never the data you have. It was your ability to ask and act across all of it.**

Axonius AI

ASK AND ACT

AI Companion

Ask anything and act on it, in a conversation.

WORK THAT RUNS ITSELF

AI Workforce

Virtual Workers watch your environment and act inside the limits you set

EXTEND


MCP Server

Bring Axonius data into the AI tools and agents you already run.

AI Companion

One agent plans, routes, and loads the right domain skill for every question. You never pick a mode or a specialist.

- **Ask:** query any asset type in plain language. Investigate a host, trace an IP, or see what changed.
- **Act, with a preview:** save queries, pin charts, tag assets, build enforcements. Every write shows what will change before anything is applied.
- **Schedule:** save any question as a recurring Scheduled Task, on a timer or after each discovery cycle.
- **Learn:** Platform tutorials run on your own live data.



TRY ASKING

"Show me all Windows servers missing our EDR agent, seen in the last 7 days."


"Why did my device count spike this week? Split the new devices by data source."

"Which users have an AI app installed but no SSO assignment?"

"Tag all devices missing both EDR tools as No EDR." (Shows a preview first.)

AI Workforce

Most enterprise AI waits for a prompt, answers, and marks the task as completed. The **Axonius AI Workforce is proactive**. It's a standing team of Virtual Workers organized by function that run continuously, track whether things are getting better or worse, and can act on their own inside boundaries you define. **This is the agentic part of Axonius AI, and the part with no real parallel in enterprise cybersecurity tools.**



Shadow to Live. Every action starts in Shadow mode, showing what a Virtual Worker would have tagged or ticketed. After a few clean cycles, flip it to Live. Standing instructions, like "ignore lab IP ranges," make each check tailored to your environment.

WHY CONTEXT IS THE BOTTLENECK


Only 45% of organizations consolidate asset and exposure data into a single view – and just 32% apply ownership consistently across cloud, EDR, and CMDB. AI can't reason well over context that's fragmented.

Source: Axonius Actionability Report 2026, conducted with the Ponemon Institute



Tip

Axonius docs are now AI-ready. Feed any page into your AI tools so that it can learn about Axonius and optimize your experience.




"The value of data is measured in connections. Axonius AI connects those dots across our environment, turning tasks that once took hours into actionable insights in seconds."

Mike Chasteen, staff security architect at Exemplar Luxury Group (formerly Saks Global)

Axonius MCP Server

The same governed asset model, available outside the platform. Connect Claude, your IDE, or custom agents and automations to Axonius data, so your AI tools get a correlated answer about an asset, its owner, and its exposure instead of a guess.



Why security teams choose Axonius AI

Benefits

Answers in seconds. Complex cross-tool analysis, answered in seconds.

- **No proprietary query language.** Every analyst gets senior-level reach on day one.
- **Fewer blind spots.** Coverage gaps, shadow AI, and stale data surface proactively.
- **AI that earns its autonomy.** Agents show their work before you go live.

Built to trust

- Runs inside your Axonius instance
- Inherits your RBAC role and data scope, with settings to further scope agent functions
- Your data stays in the tenant; zero retention at the AI layer
- Every write is previewed and confirmed - the gate is server-enforced.

Ready to try Axonius AI
on your asset data?

Learn more at
<http://axonius.com/platform/ai-for-security>