AEON · CYBERGUARD

AI agent security

Turn "how do you secure your AI agents?" from a deal-stopper into a deal-closer.

Your biggest customers' security reviews now ask how you secure your AI agents, and a weak answer stalls the deal. These systems reach into your data, call your tools, and act for users - the failure modes block deals and become incidents.

Access Can one user, tenant, or role reach another's data through your agents? (cross-customer breach)

Execute Can your agents or tools be driven to take unauthorized actions? (tampering, account takeover)

Leak Can your AI be turned into an exfiltration or SSRF path? (data loss, internal pivot)

What you get

An authorized, evidence-first AI Agent Security Review. In 10 business days: a verified risk register across access, execute, and leak; reproduction-safe evidence; developer-ready fixes; a retest to closure; and a buyer-shareable evidence package - every finding mapped to OWASP LLM Top 10, MITRE ATLAS, and your SOC 2 / NIST AI RMF controls.

Why trust us

- Authorized-only, scope-disciplined - a signed rules-of-engagement; we test only what you authorize.
- Owned-lab research pipeline in the exact class of tools AI teams adopt (anonymized, coordinated).
- Human-verified, AI-assisted - never an unreviewed machine report.
- Our wedge vs a generalist pentest or bounty: the AI-specific trust boundaries (agent authorization, tool command-boundary, MCP, injection-to-action) that generalist scopes miss.

Audit-ready evidence, mapped to your controls - your auditor still issues the opinion. A compliance platform confirms a control exists; we prove your agents cannot be abused.

Start free. A public-exposure snapshot and a 60-minute threat-model call, no access. Then a fixed-fee diagnostic or full Review, and an ongoing Release Gate retainer, under a signed ROE.

Contact Aeon

Aeon AI Risk Management · airiskmanagement.ca/contact · info@airiskmanagement.ca