Bessemer Venture Partners

The CISO action plan for securing AI agents

AI agents are transforming enterprise operations — and exposing security gaps existing tools weren't built to close. Here's how to orient your defensive strategy.

> **"The core risk isn't vulnerability, it's unbounded capability."**
>
> — Barak Turovsky, Bessemer Operating Advisor, former Chief AI Officer of General Motors

Internal audit: 7 questions to ask your team

- [ ] How extensively are AI agents deployed in your environment today?
- [ ] What's your biggest concern about their security risks?
- [ ] Do you care more about coding agents (Cursor, Claude) or generic ones?
- [ ] Which layer makes most sense for AI security controls — end point, network, or identity?
- [ ] Is ther room for purpose-built, agent-specific solutions?
- [ ] With so many AI security startups emerging, how do you distinguish between them?
- [ ] Are you more focused on visibility into what agents are doing, or on preventing them from being compromised?

5 priorities to close the protection gap

1. Align on risk posture before buying anything	Know your organization's position on agents — all in, cautious, or wait-and-see. The framework should follow the strategy, not precede it.
2. Treat agents like production infrastructure, not applications	Ownership first, then constraints, then monitoring. Define responsibility, limit permissions, and enforce guardrails before any monitoring tool is turned on.
3. Start narrow, then expand deliberately	Launch agents with minimum permissions for a specific task. Validate behavior in constrained environments before expanding access.
4. Close the freedom-versus-control gap with guardrails, not monitoring	Monitoring tells you what an agent did. Guardrails determine what it's allowed to do. Define boundaries at the action level, not just the access level.
5. Give every agent an identity — and treat it like an employee	Every agent needs a managed identity with scoped authentication. If you can't answer "what can this agent do, on whose behalf, and who approved it" — you're not ready.

The bottom line

AI agents are very likely already inside your enterprise — and the security infrastructure to govern them doesn't exist yet. The CISOs who build it deliberately now will define what safe AI deployment looks like, and the ones who don't will spend the next decade cleaning up breaches that were entirely foreseeable.