AI Agent Security Is a New Recurring Revenue Line for Your MSP.

Your clients are deploying AI agents now. Nothing in their security stack covers the agent layer. Crawdad is the managed AI-security service you deliver — a new service line layered on top of the relationships you've already built.

THE OPPORTUNITY

A new managed-service line no one else is offering yet.

THE REVENUE

Recurring, not one-time

Crawdad is a managed service, not a product sale. Layer it into your existing security practice alongside EDR, DLP, and identity. One more line of recurring revenue on every client — channel-friendly economics built for how MSPs sell.

FIRST MOVER

Define the category

AI agent security is nascent. The MSP who brings a real answer to their clients now sets the standard for how this layer gets covered. First movers define the category — everyone else follows.

EVERY CLIENT

Universal applicability

Every client deploying AI agents needs this coverage — regardless of size, industry, or maturity. Regulated clients in finance, healthcare, and government feel it first and will ask. You bring the answer.

WHY NOW

Your clients are exposed right now.

AI agents are moving from pilot to production across every industry. They act with real authority — your client's credentials, files, network access. When an agent reads a poisoned document, it follows hidden instructions silently. This is a net-new attack surface that EDR, DLP, firewalls, and identity tools don't see.

Regulated clients are already asking who secures their AI agents. The security buyer is looking for answers — and the MSP who arrives first with a real one wins the room. The gap between "my clients are deploying agents" and "I can secure them" is the opportunity.

WHY YOU

You already own the security relationship. Extend it.

Your clients trust you with their endpoints, their networks, their identity. Agent security is a natural extension. Nothing in the typical MSP stack — EDR, DLP, firewall, identity — covers the agent layer. Being the one who covers it makes you the leader. Crawdad doesn't replace anything in your current stack. It covers the surface that nothing else sees — additive to your offering, no displacement, no rip-and-replace.

HOW IT FITS

Fleet-scale agent security. Every client. One console.

- ✓ One-command deployment per endpoint
- ✓ Fleet Console: manage AI security across your entire client base
- ✓ Per-client tenancy and isolation
- ✓ Centralized policy and detection monitoring
- ✓ Scales to 500+ devices with bulk operations
- ✓ Local-first — raw client content never leaves their machines by default
- ✓ Metadata-only telemetry, customer-governed
- ✓ Fits alongside EDR, DLP, identity — no rip-and-replace

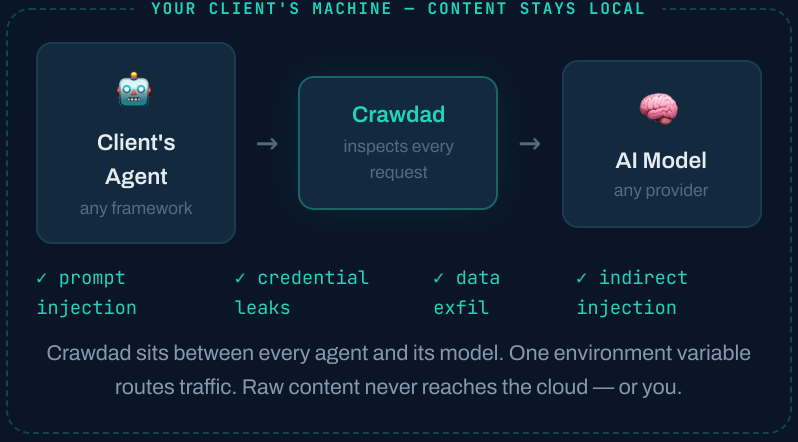

PROOF

A public, reproducible benchmark. We'd rather you verify than take our word.

99.8%
detection rate

497	1,669	0.09%	1
attacks tested	total samples	false-positive rate	missed attack

Every result published openly. No other vendor in the category has published a reproducible benchmark. Clone the corpus, run it yourself, compare your tool. Multi-layer detection pipeline, built in Rust, 2,988 automated tests across 26 crates.

Zero-knowledge architecture. Raw content never leaves the client's device by default. Metadata-only telemetry, customer-governed. When you pitch Crawdad to your clients, you can tell them their data stays on their machines — because it does. That's a conversation no other AI security vendor can have.

`Benchmark corpus: CC BY 4.0 · github.com/AndrewSispoidis/contemporary-agent-attacks · 497 attacks · 1,172 benign negatives · 22 categories`

The MSPs who secure AI agents first define the category.

Talk to us about becoming a Crawdad partner. We'll walk you through the deployment model, the economics, and how to position agent security to your clients.

contact@getcrawdad.dev

getcrawdad.dev · Crawdad Security, Inc.