



THE AGENT MANAGEMENT LAYER

Let your clients run AI agents — **without** **handing over the keys.**

AI agents are about to act on your clients' ERPs, CRMs and finance systems. Modgard is the layer in between: it gives every agent a **name**, a **rule set**, a **budget**, a **human approval step** for the risky stuff, and a **tamper-evident record** of everything it did — operated from one screen, by you, on behalf of the client.

— THE PAIN IT SOLVES

"My client wants AI agents in their systems. Who's accountable when one does something it shouldn't?"

Today that question has no good answer. An unidentified agent can call any tool on any connected system, and the log can't even say which agent did it. No MSP can put their name on that. The tools you already sell — RMM, PSA, PAM — govern **human technicians** beautifully. None of them govern the client's **AI agents**. That gap is the client's risk, and right now it's yours too.

— WHAT THE OPERATOR SEES — THE FIVE PROMISES

One screen. Named agents, rules, approvals, budgets, and a record.

Every action an agent takes routes through Modgard and resolves to one of three states:

● Allowed – acts freely, on budget

● Held – paused for your approval

● Refused – not permitted

01

Every agent is named

An agent connects and shows up by name — with what it did, on which systems. Identity is derived by Modgard, not self-reported.

Value: "which agent did this?" finally has an answer. That's the line between a product you can be liable for and one you can't.

02

Rules per system, in plain language

Each connected system reads *Asks first, Acts & notifies*, or *Acts freely* — with the tools it's allowed and the budget it has. Loosen a system only when it earns it.

Value: least-privilege by default, tuned per client, run by a normal operator — not a security specialist.

03

Risky actions pause for a human

An action above the agent's permission — HR data, a large write, a blown budget — lands in a *Waiting for you* queue with the exact request, reason and cost. Approve & run, or decline.

Value: the human-in-the-loop control that lets you say yes to agentic AI without handing a client's ERP to an unsupervised bot.

04

A budget and an off-switch

Per-system spend caps that actually hold. The meter climbs as agents act; the moment it hits the cap, the next action pauses instead of running.

Value: cost and blast-radius control, per client. No surprise bills, no runaway loops.

05

A permanent, tamper-evident record

Every ask, refusal and approval is in a cryptographically chained audit the operator can verify — it even proves *if* anything was ever altered, and that everything since is clean.

Value: the artifact you hand the client as compliance evidence. "If it isn't in the record, it didn't happen" — applied to AI.

— WHY IT'S A KILLER FEATURE — THE WHITE SPACE

- ✗ **Kaseya & ConnectWise** ship AI — but embedded in *their own* tools. As of mid-2026 neither governs the *client's* agents.
- ✗ **The nearest gateway (WYRE)** points AI at the *MSP's own* tools — allowlists and log-shipping, no approvals, no budgets, no evidence.
- ✗ **RMM / PSA / PAM** govern human technicians brilliantly — and were never built for an AI agent as the principal.
- ✓ **Modgard** governs the *client's* AI agents — and does it in the channel's own language: just-in-time access, propose-and-wait, ticket-anchored record. A familiar idiom, extended to a new kind of user.

New line item

RECURRING REVENUE

A governance service you bill per client, per month — on top of the systems you already manage for them.

You're already there

ZERO NEW RELATIONSHIP

You already hold the client's systems and their trust. This is the natural next service, not a new sale to a new buyer.

First-mover

COMPETITIVE MOAT

Be the MSP in your market who can safely say "yes, we govern your AI agents" — while competitors are still saying "we don't touch that."

- **Stickier accounts.** Governance sits under every agent the client runs — the harder it is to rip out, the longer they stay.
- **Risk you can price.** You turn "AI is scary and ungoverned" into a defined, billable service with a clear deliverable: the compliance record.
- **Sold your way.** Single-tenant, per client, operated from one screen — it fits how you already run managed services.

— THE ONE-LINE PITCH

Your clients are about to let AI agents into their most important systems. **Modgard is how you let that happen — safely, on the record, and as a service you sell.**

"We don't just connect your systems to AI. We govern it — and we prove it."