



THE CONTROL PLANE FOR MANAGED AI OPERATIONS

Run AI in your clients' businesses — without betting your company on it.

Your clients are asking for AI. Deploying it badly is an existential risk to your MSP; not deploying it hands your clients to whoever will. Modgard is the operator plane in between — **agents you deploy per client, autonomy you set, credentials you hold, an audit trail that answers the lawyer, and economics that show your margin.** It also heals the legacy integrations you're still stuck maintaining, and counts the hours it gives back.

— THE PAIN IT SOLVES

"My clients want AI in their systems. I can't say no — and I can't say yes to something I have no way to control, explain, or bill."

Right now that gets done with spreadsheets, shared credentials and hope. An unidentified agent can call any tool on any connected system, and the log can't say which agent did it — while shadow AI is already inside your clients and already your liability. The tools you sell today govern **human technicians** beautifully: RMM, PSA and PAM were built for people with badges. **None of them govern software that acts on its own.** Meanwhile the legacy integrations still break every week, and that toil is what's eating the margin you'd fund the AI service line with.

— ONE PLANE, TWO WORKLOADS

The wedge pays today. The bet compounds.

Both run on the same control plane, the same audit, the same per-client economics — so you sell one platform, not two.

BROWNFIELD · THE REVENUE WEDGE

Heal the integrations you already babysit

Modgard detects schema and contract drift, re-enters the mapping loop on its own, trips a circuit breaker when a connector misbehaves, and reports the ticket-minutes it displaced. Sixteen vendor dialect packs — SAP IDoc and Infor ION BOD among them.

Priced against KTLO labour displaced — it pays for itself on day one.

AGENTIC · THE LAND GRAB

Govern the agents arriving next

Every agent named and scoped, tool grants enforced against a blueprint, hard spend caps, a three-layer kill switch, and human approval on the risky calls — all on the live tool-call path, refused *before* the client's system is touched.

The accountability layer insurers, auditors and clients are about to demand.

— WHAT THE OPERATOR SEES — THE FIVE PROMISES

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

Every action — by an agent or by you — routes through Modgard and resolves to one of three states:

● Allowed — acts freely, on budget

● Held — waiting for a human

● Refused — and it names the rule

01 Deploy per client, from one console

Stand up a client's engine in minutes — provisioned into real cloud infrastructure with live step-by-step progress, then reachable from your console through a secure relay. One multi-tenant plane for you, a single-tenant engine per client.

Value: onboarding is a workflow, not a project — so a new client is billable this week, not next quarter.

02 Autonomy you dial, per client and per system

Each connected system reads *asks first*, *acts & notifies*, or *acts freely* — with the tools it's allowed and the budget it holds. Anything outside the blueprint is refused and raised as a security event. Loosen a system only when it earns it — and Modgard suggests promotions once behaviour has been clean, but never takes them itself.

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

03 A human gates the risky stuff — including the AI's own ideas

Risky actions land in a *waiting on you* queue with the exact request, reason and cost. Raven, the built-in copilot, can investigate the fleet and *propose* a fix — but proposals are all it can make.

Approval is a human decision, gated by role, and it's the approval that reaches customer infrastructure.

Value: you say yes to agentic AI without handing a client's ERP to an unsupervised bot — and the AI operates under the same rules as everyone else.

04

Fleet-scale in one decision

When a finding hits many clients — an out-of-date engine across the estate — Modgard gathers exactly the affected set and stages one batch. A single approval fans it out per-tenant and reports back honestly: *patched 6 of 7; tenant 5 is offline, queued, applies on reconnect*.

Value: more clients per technician. The work stops scaling with headcount, and that difference is your margin.

05

A permanent, tamper-evident record

Every ask, refusal, approval and change is written to a cryptographically chained audit — recording not just what happened but *why* it was done. Export it as an evidence pack with its own verification report; it proves whether anything was ever altered, and that everything since is clean.

Value: the artifact you hand the client, the auditor, or the insurer. "If it isn't in the record, it didn't happen" — applied to AI.

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

- ✗ **RMM & PSA incumbents** govern human technicians, and will bolt on AI modules eventually — but they are built around a badge and a ticket, not an actor that acts on its own.
- ✗ **Agent infrastructure vendors** arm the people *building* agents. They do not serve the operator running an agent estate inside somebody else's business — they're your supply chain, not your competition.
- ✗ **Generic iPaaS** connects systems for a single org — no multi-tenant fleet, no per-client autonomy policy, no channel economics, no evidence pack.
- ✓ **Modgard** is the operator plane above all of them — multi-tenant, deployed into other people's businesses, with autonomy dials, held credentials, a verifiable audit chain and per-client margin math. Built for the person who is *accountable*, not the person who is building.

— WHY AN MSP RESELLS IT

Free plane

RECURRING REVENUE

The console you get addicted to is never taxed. Revenue is per-client wholesale plus outcome share — so adding the tenth client is a margin event, not a licence negotiation.

A new service line

MANAGED AI, SOLD BY YOU

"We govern the AI in your business" is a service your clients are already trying to buy, and nobody in your market is selling it yet. Land it before the RMM vendors ship a module.

Proof on demand

TRUST YOU CAN BILL

The exportable, verifiable record is your QBR evidence and your SOC-2 / insurance pack — a deliverable your competitors literally cannot hand over.

- **It pays for itself first.** Integration healing displaces KTLO hours from day one — that's what funds the AI service line, rather than betting the P&L on it.
- **PSA citizenship.** ConnectWise, Autotask and Halo — bidirectional, so closing the ticket closes the incident. If it doesn't live in your ticket flow, it doesn't exist in your workflow.
- **Stickier than anything you sell.** Once you hold the credentials, set the autonomy and produce the audit, you are woven through every AI decision in that client's business.
- **It tells you the truth.** When something needs a human, Modgard says which kind — a judgment call that should stay yours, or something not yet automated. It never dresses a gap up as a governance principle.

— THE ONE-LINE PITCH

Your clients are going to run AI in their businesses with or without you. **Modgard is how you run it for them — with autonomy you control, credentials you hold, and an audit trail that answers the lawyer.**

"Someone is going to be accountable for the AI in your business. We're built so that it can be us."

MODGARD · THE CONTROL PLANE · **deploy** · **govern** · **prove**

Multi-tenant console · per-client engine · operated by **you**