Agent governance, the short version

What ClawMetry is, who inside your organisation it is for, what it costs, and how a rollout actually goes. Written to be forwarded.

The problem you already have

Your engineers adopted AI coding agents without waiting for a procurement cycle. That was the right call for velocity and it left three questions unanswered:

- **What is this costing?** Agent spend arrives as one API bill with no attribution to teams, repositories, or people. Nobody can tell a productive $40 session from a runaway loop that burned $40 achieving nothing.
- **What are they allowed to do?** An agent that can run shell commands and open pull requests has real blast radius. Most organisations have no policy layer between an agent's intent and its action.
- **What did they actually do?** When someone asks what touched a system six weeks ago, "it was an agent" is not an answer that survives an audit.

What ClawMetry does

See	Cost and activity attributed per agent, session, model, repository, tool, and person. Waste and runaway loops flagged against your own baseline rather than a generic threshold.
Control	Budget ceilings, call-level tool policy, and an approval queue that holds a risky action for a human decision instead of logging it afterwards.
Prove	A hash-chained audit record, exportable to your SIEM, plus generated evidence bundles mapped to NIST AI RMF and SOC 2 control families.

Coverage

Twenty-six agent runtimes, including Claude Code, Codex, Cursor, Copilot, Gemini CLI, Cline, OpenHands, Aider, Goose, Devin, and self-hosted harnesses. Your engineers did not standardise on one agent; the control plane should not require them to.

Enforcement depth varies by runtime, because some expose a pre-action hook we can gate on and others only emit telemetry after the fact. We publish a per-runtime conformance matrix and will walk you through where your specific runtimes land before you buy.

Who this is for

Platform engineering	Owns the rollout. Wants attribution, a fleet view, and limits that do not create support tickets.
Security	Owns the risk. Wants to know what agents can reach, an approval path for the dangerous actions, and an audit trail in the SIEM.
Engineering leadership	Owns the budget. Wants a defensible number for what agents cost and evidence that the spend correlates with output.

Deals usually start with the first, get blocked or unblocked by the second, and are signed by the third.

Deployment

Collection is local in every mode — a daemon on the same machine as the agent, reading what the agent already writes.

MODE	WHO OPERATES THE DASHBOARD	WHAT LEAVES YOUR NETWORK
Cloud	Rangehead	A snapshot encrypted on your machine with a key we do not hold
Self-hosted	You	Nothing
Air-gapped	You	Nothing, and no licence callback either

The daemon opens no inbound ports. It requires Python 3.9 or newer and runs on macOS, Linux, and Windows.

What a rollout looks like

Weeks 1–2 · Observe	One team, read-only. No behaviour change for engineers. You get a baseline of real cost and activity — which is where most organisations find their first surprise.
Weeks 3–6 · Widen	The rest of engineering. Budget alerts on, limits off, so thresholds get set from your data instead of a guess.
Weeks 7–12 · Govern	Enable the controls you now have evidence for. Wire the audit stream to your SIEM. Generate the first evidence bundle.

We resist reversing that order. Enforcement configured before a baseline exists is enforcement configured on a guess, and the first false block costs more trust than the control was worth.

Commercials

Enterprise engagements start at $25,000 a year, priced on the number of machines running agents, the deployment mode, and the support commitment. Annual invoicing.

Smaller teams are better served by the self-serve plans at clawmetry.com, and we will tell you if that is you rather than sell you a contract you do not need.

Why a small vendor is a defensible choice here

The honest version: we are a small company, we are not SOC 2 certified, and single sign-on is still in development. Those are real considerations and they are documented rather than buried.

Against that: ClawMetry is MIT-licensed and public, so your security team can read the code before procurement starts — which is more than most certified vendors will offer. It has more than 260,000 installs. And the deployment modes mean you can adopt it without any of your data reaching us at all.

THE NEXT STEP

Ask for the security and architecture whitepaper and the procurement pack — both are written and ready. Most evaluations get most of the way through on documents before anyone books a meeting, and we would rather you did it that way.

vivek@rangehead.com · rangehead.com/contact

Rangehead is a trading name of Instalabs LLC, Wyoming, United States.
Rangehead is the commercial home of ClawMetry, an open-source project at github.com/vivekchand/clawmetry.
Questions: vivek@rangehead.com · rangehead.com
This document describes the product as at the version and date above. Where it conflicts with a signed agreement, the agreement governs.