11:59
ready for tomorrow

COMMERCIAL

WHERE ALL THREE LEAD · 11:59 OPERATE

AI that keeps working after launch.

You have agents a contractor built, tools your staff stood up, and no team to keep any of it running. It worked on day one. Now the answers are a little worse, the costs a little higher, and nobody can say what it delivers.

Trusted by **TaxAudit, Horrocks, Battelle Energy Alliance, Navy Mutual Aid Association** and **Allmoxy** · **89+ NPS, measured every year**

11:59 Operate runs your AI agents and solutions in your cloud, whoever built them. It commits in writing to what's covered, how fast we respond, and how value is measured, and sends a monthly report leadership can read in five minutes. One predictable line item instead of a new department. You own the deployment, the data and the option to run it yourselves.

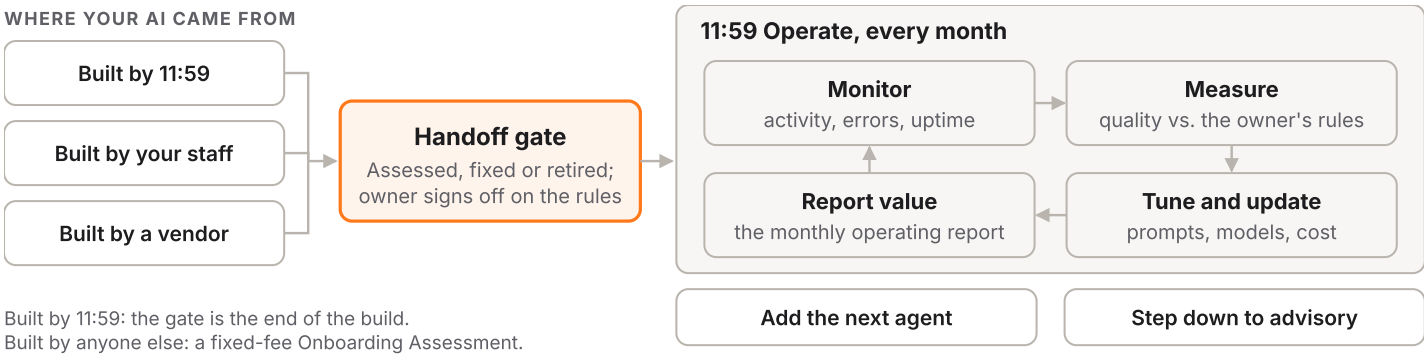
WHAT WE RUN

Five operating disciplines

DISCIPLINE	WHAT WE DO	WHAT YOU SEE
Governance	Keep each agent inside its owner's rules	An owner, autonomy level and audit trail per agent
Observability	Monitor activity and uptime; respond to incidents	A live view of each agent, and incident reports
Performance	Test on real cases; tune prompts, retrieval, models	Quality that holds steady as models change
Cost	Track spend per agent; route work to the right model	Spend per agent, and where it's coming down
Value	Measure what each agent handled and gave back	A monthly report read in five minutes

HOW IT WORKS

One gate, then a monthly loop



Built by 11:59: the gate is the end of the build.
Built by anyone else: a fixed-fee Onboarding Assessment.

EVERY MONTH

The operating report

- **The month in five lines.** items handled, staff hours returned, run cost, agents meeting their owner's bar
- **Every agent, one line.** owner, autonomy level, quality, volume, incidents, cost, hours returned
- **Decisions for you.** autonomy changes we recommend, with evidence; nothing changes without the owner's sign-off

How value is measured. Quality is the share of an agent's work that passes its owner's checks in weekly reviews of real cases. Hours returned are items handled times the minutes per item the owner agrees. We report time given back, not cash savings.

BUYING

Tiers

TIER	COVERS	INCLUDES
Foundation	A small set of agents or solutions	The five disciplines, business-hours support, monthly report, named operations lead
Growth	Agents across several departments	Plus a monthly cost pass and new agents from the menu
Enterprise	A larger estate and the Agent Foundation	Plus round-the-clock response to the most severe incidents and audit packages
Advisory	Your team runs the agents	Monthly quality and governance review, quarterly audit, on-call escalation

Term and exit. Each tier is a flat monthly fee on an annual term. If you leave, we hand over every runbook, test, prompt and configuration, and the agents keep running in your cloud. An agent acting outside its owner's rules is always the most severe incident: paused first, investigated second.

WHY 11:59

An AI team you don't have to hire

- **We run what we build.** The same forward deployed specialists who built your agents operate them.
- **Your cloud, your data.** Everything runs in your AWS or Azure environment, under your identity provider.
- **A capability, not a dependency.** Runbooks, tests and documentation are yours; Operate steps down as your team steps up.

> *AI that keeps working after launch, for one predictable monthly fee.*

Next step: an Onboarding Assessment or an AI Rapid Scan

Already have agents running? An Onboarding Assessment gives a verdict on every agent, a remediation plan and a recommended tier. Starting fresh? Start with an AI Rapid Scan. Both are scoped and quoted before we begin.

Fixed scope and fee, quoted before you start. 11:59 services are also listed on **AWS Marketplace**, and AWS funding programs may offset part of the cost.

11:59 · Sacramento, California
hello@1159.ai · 888-369-1159
www.1159.ai