11:59
ready for tomorrow

COMMERCIAL

WHERE ALL THREE LEAD · 11:59 OPERATE

# 11:59 Operate

## AI that keeps working after launch.

**We run it. You own it.**

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.

11:59 Operate is the managed service that runs it. We take on agents, solutions and agent layers wherever they came from, operate them in your cloud under your identity and audit rules, and commit in writing to what's covered, how fast we respond, and how value is measured.

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

> *The deployment, the data and the option to run it yourselves always stay with you.*

THE PROBLEM

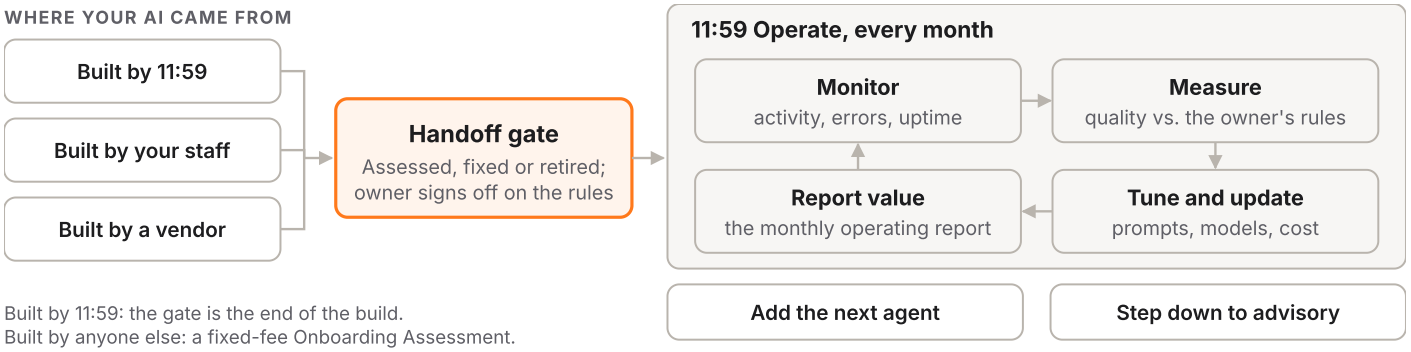
## Why AI quietly stalls after launch

Most AI work is built, handed over and left. It works on day one. Months later the answers are a little worse, the costs are a little higher, and nobody can say what it's delivering.

| | BUILT AND HANDED OVER | OPERATED BY 11:59 |
|---|---|---|
| **Quality** | Drifts as models, prompts and data change | Measured every week against the owner's rules, tuned when it slips |
| **Governance** | A policy written at launch | Enforced in the work: every agent has an owner, an autonomy level and an audit trail |
| **Cost** | Model and platform spend grows unwatched | Usage and model costs tracked per agent and optimized every month |
| **Value** | Nobody reports on it after launch | A monthly report on what each agent did and what it gave back |
| **Next agent** | A new project, a new procurement | Added from a menu, on the same foundation |
| **Result** | Risk you can't see | AI that keeps paying for itself |

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.

Wherever an agent came from, it passes the handoff gate before Operate takes it on. From then on it runs on the same monthly loop, new agents join the same loop, and Operate steps down to advisory when your team is ready to run it.

WHAT WE RUN

## Five operating disciplines

Every agent, solution and agent layer under 11:59 Operate gets the same five.

| DISCIPLINE | WHAT WE DO | WHAT YOU SEE |
|---|---|---|
| **Governance** | Keep each agent inside the rules its owner set: access, autonomy level, data boundaries. Prepare audit packages on request. | An owner and an autonomy level for every agent, and an audit trail you can hand to an auditor |
| **Observability** | Monitor every agent's activity, errors and uptime, and respond to incidents | A live view of what each agent is doing, and incident reports when something goes wrong |
| **Performance** | Test answers and actions against real cases, tune prompts and retrieval, manage model updates | Quality scores per agent that hold steady or improve as models change |
| **Cost** | Track model and platform spend per agent; route work to the right model; refresh the License Posture report every quarter | Spend per agent, and where it's coming down |
| **Value** | Measure what each agent handled and the time it gave back | A monthly operating report your leadership can read in five minutes |

As the estate grows, we add new agents and workflows from a fixed-fee menu, on the same foundation, so the next one never starts from zero.

GETTING STARTED

## Two ways in

| | BUILT BY 11:59 | ALREADY RUNNING, BUILT BY SOMEONE ELSE |
|---|---|---|
| **Who it's for** | Clients coming from Standardized AI Adoption, the AI Solutions Catalog or Reduce Your SaaS Costs | Organizations with agents their staff, a vendor or a contractor built |

| | BUILT BY 11:59 | ALREADY RUNNING, BUILT BY SOMEONE ELSE |
|---|---|---|
| **First step** | The handoff gate at the end of the build | A fixed-fee Onboarding Assessment of the agents you have today |
| **What the assessment covers** | Already done during the build | What each agent does, who owns it, what it can reach, how it's tested, what it costs, and what must be fixed first |
| **Then** | Operate takes it on | Agents that pass join Operate; the rest get a remediation plan, or a recommendation to retire them |

Nothing joins Operate until its owner has signed an operating acceptance covering what it does, what "healthy" means and its autonomy level.

ROLES

## Who owns what

You own the decisions and the assets. We do the running.

| | YOUR TEAM | 11:59 |
|---|---|---|
| **What each agent should do** | The owner sets the rules and what "healthy" means | Turns those rules into tests and enforces them |
| **Autonomy** | Approves every move up the ladder: human-led, assisted, autonomous | Recommends when an agent is ready, with the evidence |
| **Cloud, data and identity** | Owns the tenant, the data and who can access what | Operates inside it, under your controls |
| **Day-to-day operation** | Uses the agents | Monitors, tunes, updates, supports and reports |
| **New agents** | Picks priorities with leadership | Builds and adds them on the shared foundation |
| **Exit** | Can take over at any time | Hands over runbooks, tests and documentation |

Most AI services firms build a dependency. We build a capability: as your people learn to run more of it, Operate steps down to advisory.

EVERY MONTH

## The operating report

- **The month in five lines.** items handled, staff hours returned, run cost, and how many agents meet their owner's quality bar
- **Every agent, one line each.** owner, autonomy level, quality, volume, incidents, cost and hours returned
- **Service levels.** each commitment, its target, and what happened
- **What happened and what we did.** every incident, fix and cost change, in plain language
- **Decisions for you.** autonomy changes we recommend, with evidence; nothing changes without the owner's sign-off
- **Coming next month.** new agents joining and the next License Posture refresh

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 equal items handled multiplied by the minutes per item each owner agrees at the handoff gate. We report time given back to staff, not cash savings. A sample report is available on request.

SERVICE LEVELS

## What we commit to

Availability and response targets are set per contract, by tier.

| COMMITMENT | TARGET |
|---|---|
| **Platform availability** | Set in your contract |
| **Most severe incidents (agent down or acting outside its rules)** | Set in your contract |
| **Degraded agent** | Set in your contract |
| **Minor issue** | Set in your contract |
| **Quality review of every agent** | Weekly |
| **Monthly operating report** | By the fifth business day of each month |

An agent that acts outside its owner's rules is always the most severe incident. We pause it first and investigate second.

BUYING

## Tiers

| TIER | WHAT'S COVERED | 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 | Everything in Foundation, plus a monthly cost optimization pass and new agents from the menu |
| **Enterprise** | A larger estate, the Agent Foundation and quarterly License Posture reports | Everything in Growth, plus round-the-clock response to the most severe incidents and audit packages on request |

| TIER | WHAT'S COVERED | INCLUDES |
|---|---|---|
| **Advisory** | Your team runs the agents | Monthly quality and governance review, quarterly audit, on-call escalation |
| **Onboarding Assessment** | Agents you already have, built by anyone | A fixed-fee review of each agent, a handoff-gate verdict and a remediation plan |

Each tier is a flat monthly fee on an annual term. New agents beyond a tier's allowance are added from a fixed-fee menu. If you leave, or move to Advisory, we hand over every runbook, test, prompt and configuration, and the agents keep running in your cloud. There's no fee to take back what's already yours.

FOR COMMERCIAL LEADERS

## An AI team you don't have to hire

Operate gives you a predictable line item instead of a new department. We assess what you have, take on what's sound, fix or retire the rest, and report every month on what it delivers.

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

WHY 11:59

## Built for the long run

- **Safe to scale.** Every agent has an owner, tests and an audit trail, and runs in your own AWS account. AWS Advanced Tier Partner with the AWS Generative AI Competency.
- **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.
- **A record of delivery.** Every AI and cloud project delivered since 2022, and 89+ NPS, measured every year.

### 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, and Operate takes over at the handoff gate when your first agents go live. Both are scoped and quoted before we begin.

Scoped and 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