



FOR THE CFO & FINANCE TEAM

## Enterprise AI That *Pays Back.*

### The Spend *Question*

AI budgets are under more scrutiny than almost any line on your forecast — and most enterprise AI pilots still can't show a return. The spend is real and immediate; the value is vague and deferred. That gap is where AI initiatives quietly die in the budget review.

ClearData AI is built to close it. It runs as a fully managed service, or deploys into **your own Azure tenant**, and **instruments every agent run** — tokens, time, tool calls, and outcomes are logged to your data lake. Value isn't asserted on a slide; it is *observed in your own telemetry*. Below is the exact model we use to build — and defend — the number, line by line.

#### HOW WE CALCULATE VALUE

**Monthly Value Created = Labor Savings + Revenue Acceleration + Error & Compliance Savings**

where  $Labor Savings = FTE\ Hours\ Automated \times Blended\ Cost / Hr \times Productivity\ Multiplier$

$Net\ Value = Value\ Created - Total\ Investment$

$ROI = (Value\ Created \div Total\ Investment) - 1$

#### THE NUMBERS, CONSERVATIVELY

**2–6 mo**

CONSERVATIVE PAYBACK

**2–5×**

NET VALUE VS.  
MONTHLY COST

**\$5K–\$90K**

NET VALUE / MONTH  
RANGE

**100%**

VALUE SEEN IN YOUR  
TELEMETRY



## The Two Sides of the *Ledger*

A ClearData AI deployment has a transparent cost structure and three measurable sources of value. You fund a platform and the **agent cluster** that runs on it; in return, work that used to consume people, time, and error-prone manual steps gets automated, accelerated, and de-risked.

### What You Invest


- **Platform subscription.** Tiered access to the SuperIntelligence platform — the dominant line item, sized to your tier.
- **The agent cluster.** Fleets, orchestrators, sub-agents, MCP integrations, copilots, and middleware — the digital workforce that does the work. Bundle grows progressively with tier.
- **Seats, runs & schedules.** Nominally priced (sliding rates per tier) so growth is friction-free. Hard caps enforce the upgrade conversation cleanly.
- **Bundled SLA — included.** Each tier ships with a matching SLA — Basic at no extra cost, Sovereign for regulated workloads.
- **Model usage — your keys.** Bring your own LLM keys (Claude, GPT, Gemini, Azure OpenAI). You pay providers directly — no markup.
- **One-time onboarding.** Setup + training, sized to tier. Smoothed jumps — no surprise cliffs between Standard and Premium.

### What It Creates

- ✓ **Labor cost savings.** Hours of skilled work — research, drafting, analysis, data ops — automated and redeployed.
- ✓ **Revenue acceleration.** Faster cycle times and more pipeline touched, capacity unlocked without new headcount. Upside, not the basis of the case.
- ✓ **Error & compliance savings.** Fewer manual errors, less rework, complete audit trail on every action — lowering the cost of getting it wrong.
- ✓ **Capacity that compounds.** Agents run in parallel and around the clock. Throughput scales with configuration, not with hiring.

### THE LEDGER SCALES WITH YOU

Monthly recurring cost (navy) against the value it creates (teal), conservative case — the gain widens as the deployment grows.





## Deploy It *Your Way*

How you consume ClearData AI shapes your cost structure. Run it as a fully managed **SaaS** subscription, or deploy the platform into your own Azure tenant as a premium **PaaS** engagement — then, for high-volume work, take the cost down further by bringing your own compute.

### SaaS — *Fully Managed*

*Fastest to value · pure operating expense*

- **We host and operate it.** No infrastructure to provision, patch, or scale — the platform is run for you.
- **Pure opex.** Predictable monthly subscription from \$2,500 (Basic) to \$30,000 (Sovereign). Nothing to capitalize.
- **Elastic with hard caps.** Add seats, runs, slots at sliding rates that drop as you grow — hard caps force the upgrade conversation when capacity is reached.
- **Best for** fast starts, smaller teams, and proving value before a deeper commitment.

### PaaS — *In Your Azure Tenant*

*Own the platform, flat fee · services scale to need*

- ✓ **Your tenant, your region.** Data and workloads stay inside the perimeter you already secure and govern.
- ✓ **Own the platform — flat fee.** One-time platform license (~\$120K-\$400K by tier) instead of renting it forever.
- ✓ **Services & add-ons recur, on need.** SLA, support, continuous updates, professional services, and model usage are recurring — you pay only for what you use.
- ✓ **Use your Azure commit.** Platform compute draws on the enterprise agreement and reserved capacity you've already negotiated.

## Optional: Bring Your Own *Local Compute*

On PaaS, route high-volume routine workloads to open-weight models running on compute you own — reserving frontier APIs (Claude, GPT, Gemini) for the hardest tasks. You trade a per-token operating cost for a fixed asset that pays for itself, and your data never leaves your hardware.

Per-token API spend displaced	≈ \$12,000 / mo
Local GPU compute — one-time capex	≈ \$60,000
Power + operations	≈ \$2,000 / mo

<b>Net monthly saving</b>	<b>≈ \$10,000 / mo</b>
<b>Payback period</b>	<b>≈ 6 months</b>
<b>Ongoing savings thereafter</b>	<b>≈ \$120,000 / yr</b>

*Illustrative, for a high-volume enterprise workload and sized to your token mix during scoping. Best suited to routine inference — summarization, classification, drafting — while frontier reasoning stays on the APIs.*



## SaaS Packages

Five managed packages scale with your team. Each is a predictable monthly subscription bundling the platform, a progressive agent cluster, seats with hard caps, and a matching SLA. Per-item rates *drop as you upgrade* — the more you commit, the cheaper each unit. Prefer to *own* the platform in your tenant? PaaS is a one-time license (\$120K–\$400K by tier) plus services that recur only as you need them.

SAAS PACKAGE	BASIC	STANDARD	PREMIUM	ENTERPRISE	CUSTOM
Best for	<i>a pilot</i>	<i>a team</i>	<i>mid-market</i>	<i>org-wide</i>	<i>regulated / sovereign</i>
Named user seats	5	15	35	75	100
Seat cap (upgrade trigger)	<b>10</b>	<b>30</b>	<b>75</b>	<b>200</b>	<b>unlimited</b>
Concurrent runs	3	8	20	40	50
Scheduled (autonomous) runs	2	5	12	25	40
Agent cluster (paid fleets)	1	2	4	6	7
Bundled SLA	Basic (incl.)	Standard	Premium	Enterprise	Sovereign
Volume discount	—	5%	8%	12%	10%
<b>Representative all-in / mo</b>	<b>≈ \$2,500</b>	<b>≈ \$5,500</b>	<b>≈ \$12,000</b>	<b>≈ \$22,000</b>	<b>≈ \$30,000</b>

### Included in Every Engagement: *We Keep You Current*

The frontier moves every month — and so does your platform. For as long as ClearData AI is in business, we continuously update your system with new models, capabilities, and agents, backed by your **SLA**. You buy a living platform, not a snapshot — your investment stays on the frontier instead of depreciating into a re-platforming project two years from now.



## The ROI Model at a *Glance*

The model scales with the size of the deployment. Each profile below is led by the **conservative case** — 60% adoption haircut from day one, no revenue acceleration counted at the floor, productivity reset to 1.0. Figures are illustrative and calibrated to your real baseline during scoping.

ILLUSTRATIVE · AFTER DISCOUNT	STARTUP	SMB	MID-MKT	ENTERPRISE
SaaS tier	<i>Basic</i>	<i>Standard</i>	<i>Premium</i>	<i>Enterprise</i>
Typical consumption model	<i>SaaS</i>	<i>SaaS</i>	<i>SaaS or PaaS</i>	<i>PaaS</i>
Typical profile	<i>1 team · &lt;25 staff</i>	<i>1-3 teams · 25-100</i>	<i>3-10 teams · 100-500</i>	<i>org-wide · 500+</i>
FTE hours automated / mo	80	240	800	2,000
Blended FTE cost / hr	\$55	\$65	\$80	\$95
<b>Recurring SaaS / mo</b>	<b>≈ \$2,500</b>	<b>≈ \$5,500</b>	<b>≈ \$12,000</b>	<b>≈ \$22,000</b>
<b>One-time setup + training</b>	<b>≈ \$10,000</b>	<b>≈ \$15,000</b>	<b>≈ \$37,000</b>	<b>≈ \$75,000</b>
Conservative value / mo	≈ \$7,000	≈ \$16,000	≈ \$40,000	≈ \$95,000
<b>Net value / mo (conservative)</b>	<b>≈ \$4,500</b>	<b>≈ \$10,500</b>	<b>≈ \$28,000</b>	<b>≈ \$73,000</b>
<b>Payback (conservative)</b>	<b>≈ 2 mo</b>	<b>≈ 2 mo</b>	<b>≈ 2 mo</b>	<b>≈ 2 mo</b>
ROI at default (Y1)	≈ 170%	≈ 200%	≈ 240%	≈ 330%

### Custom Tier — *Sovereign / Regulated*

For public-sector, healthcare, financial services, or sovereign-residency workloads. **≈ \$30,000 / mo recurring** with full Enterprise-grade catalog (all MCP integrations, all middleware, all copilots), dedicated environment, audit/compliance support, and Sovereign SLA (1hr 24x7 + compliance liaison). Per-item rates priced at Standard-tier — the Custom premium lives in the platform fee, not in catalog markup. Bundled SLA included.

*Conservative case = 60% adoption haircut, revenue acceleration excluded, labor capture cut to 50%, productivity multiplier reset to 1.0, and error savings halved. It is the floor we stand behind; full-model upside (no haircuts) reaches 3-5x the conservative net at each tier.*



## Built to Survive *Scrutiny*

We don't ask you to take the return on faith. We instrument it, baseline it, and count only what you'd defend in a board meeting.

### How We Defend the Number

- ✓ **Baseline first.** We measure today's cost and cycle time before go-live, so savings are deltas against your real numbers — not vendor benchmarks.
- ✓ **Meter every run.** Per-interaction telemetry — tokens, latency, tool calls, outcomes — is logged to your data lake. Value is observed, not estimated.
- ✓ **Count only what's realized.** Hours redeployed and rework avoided come first; revenue acceleration is carried as upside, never the basis of the case.
- ✓ **60% adoption haircut by default.** Day-one numbers assume 60% utilization, not 100% — the stress-test case is the floor we stand behind.

### How We De-Risk the Spend

- ✓ **Start small, expand on proof.** Begin at a tier, prove value on your own telemetry, then scale the agent cluster.
- ✓ **Hard caps force the right conversation.** Seat / run / slot caps trigger upgrade dialogue at the right moment — no shadow-IT cost creep.
- ✓ **Bring your own model keys.** LLM spend goes straight to the provider, transparent and unmarked-up by us.
- ✓ **Bounded autonomy.** Hard time and token budgets cap every autonomous run; human-in-the-loop gates consequential actions.
- ✓ **No lock-in.** Your configuration, data, and audit trail live in your tenant — portable, and yours.
- ✓ **Never goes stale.** We update the platform continuously, for as long as we're in business — under SLA, your investment stays on the frontier.

## PROVABLE VALUE. *MEASURED, NOT ASSUMED.*

ClearData AI runs in your tenant, instruments every agent run, and builds the business case on your baseline — so the return survives the budget review.

Send us your baseline. We'll build the ROI model on *your numbers*, line by line.

*Figures are illustrative planning defaults from the ClearData AI investment model, calibrated to each client's baseline during scoping. They are estimates for planning, not a guarantee of results.*