AMTECH.

AMTECH AI, Dalton, Pennsylvania. How the AI employee is built. Benjamin Palaskas, founder.

AMTECH is a platform for creating, running and managing AI employees. Each employee is a persistent team of AI agents with a computer of its own, doing a business's real work. This page says what is under it, what I built myself, and what the experiments so far have measured.

1. The employee is a team of agents on its own computer. The owner talks to one agent, the lead, which knows everything running and starts each job in its own working session, several at once. Each session has a browser, a terminal, code, files, computer use and vision, memory that survives between jobs, scheduled routines, and helper agents it can split work across. They keep working after the owner closes the app. Underneath is Hermes, Nous Research's open-source agent runtime, under the MIT licence; everything above it is AMTECH's. I fork it at v0.22.2 and own it from there; the employee image changes no line of its core today, so the fork starts identical.

2. The brain standard. Every employee carries the business's own brain: its prices, its customers, its rules, its past jobs and every decision, as plain files the owner can read and correct. The employee keeps the brain current itself, so one business's employee gets better every week instead of starting over. The same standard runs AMTECH's own business, and has since August, with about 600 recorded lessons and every job on the record. The brain is the part a competitor cannot copy by renting the same model: it is the business, written down in a form a machine can work from.

3. Agency inside, authority outside. The agents work freely inside their own computer. Anything that leaves it waits for a person the business has authorized, and who the employee is, who may approve and who may spend are enforced by code the agents cannot argue with. Secrets never enter the model's context; secret-backed work goes through deterministic adapters that return redacted receipts. The owner approves anything that reaches a customer or touches money, and every action is written down for the owner to check. Payments run on Stripe Connect with AMTECH's platform fee riding on each charge; invoicing finished work, progress invoices, card and bank payment and automatic reminders are proven end to end in Stripe's test mode.

4. The gateway. Every model call, picture and second of video any employee makes goes through AMTECH's own gateway. It meters cost and price per call, sends each call to the best model for the job from any lab, and makes each pricing experiment measurable from day one. When any lab ships a better or cheaper model, every employee gets it the same day. The labs' race is my supply chain.

5. The factory and the pool. Employees are built from one standard and kept ready, so a person at any door is handed a live employee at the first word, and the employee learns the business from its first job and writes it to its own records. No person is needed to make one; the $1,200 is the managed tier, where I work with the business myself. Next, one business runs several employees that work together, and an employee is trusted to move money inside limits the owner sets, with keys it can never read and a kill switch.

6. The public employee. The same employee can be a web page. On amtechai.com/contractors a visitor describes a job, the employee prices it from that business's own prices and starts work, answering in about a second once it is open, makes a page the customer can open, and can take payment by card or bank. Every visitor is a free trial, and a page like it can be put up for any trade. One of those runs is published whole on the site as receipt number one: a gutter job priced at $2,912 from the visitor's own rate, with its two assumptions named.

What the experiments have measured so far.

- Cost to serve, from the gateway's ledger (7,428 calls to 2026-10-17): typical employee work costs about 5 cents an hour in model calls; the busiest continuous stretch on record cost 46 cents an hour. A whole onboarding interview costs a few cents. A one-week free trial costs about $4 typical, $12 worst case.
- Delivery speed, from the commit record: a moving company's site, lead handling and back-office integrations were migrated in 4 hours 45 minutes of commits; a service company's payments and admin took about 25 hours. These are delivery times, not claims of autonomy.
- Pricing, in discovery by design: a setup fee plus usage at about an 85% margin on cost, a flat 50 cents per working hour, and revenue share (8% of a service company's revenue; 10% of a festival's money raised) run side by side, each measured by the gateway, so the data and not a guess picks the one owners keep.
- Email, from a 114-message outreach corpus: the thresholds for what reads as a person writing (paragraph count, sentence length, no tracking, no hard wrap) were measured from the corpus and are enforced as gates in code before anything sends.

Where it stands. The labs serve companies by sending engineers in, aimed at mid-sized firms; Meta and xAI launched agents for businesses this fall. AMTECH runs a team of agents for each business as a managed service, on the best model from any lab, and can be paid from what the business earns. For a service business turning over $1 million to $10 million a year, a well-configured team of agents can run essentially everything except the physical work. The 36.2 million small businesses in America hold almost half of its private-sector jobs, and an engineer per company will never reach them; a product can.

ben@amtechai.com · 805-886-9173 · amtechai.com