

EmAGI

Your next hire isn't human.

A **teammate**, not a tool. EmAGI turns hiring into deployment: a digital employee that lives in your messenger, remembers everything and does the work.

45

AGENTS IN
PRODUCTION

19

BUSINESSES

10+

PAYING CUSTOMERS

\$350K

PRE-SEED ROUND



Business still runs on people doing work by hand.

Urgent, expensive, and getting worse every year.

\$56K+ a year

One remote specialist: \$4,700 salary, \$1,100 benefits, \$900 overhead per month. For most companies that is the single largest controllable cost.

25–35% leave

Every year a quarter of them quit. Trained, then gone, and the context goes with them.

73% can't hire

Of small and mid-sized businesses cannot find the specialists they need at all. [SHRM, 2025](#)

AI tools don't close it

Chat assistants answer, they don't execute. Someone still has to operate them by hand, prompt by prompt, and they forget everything between sessions.

Who feels it

From owner-run teams to large organisations: construction, agencies, e-commerce, professional services, crypto and fintech. They need a specialist's output, not another dashboard.

The work must be done by hand, the specialist is missing or costly, and **the knowledge walks out the door with the person.**

A new category is forming right now.

TAM

\$82B

AI agent orchestration by 2035, growing from \$13.6B in 2026 at 22% CAGR. MarketsandMarkets, 2025

SAM

\$4.2B

AI employees — a conservative estimate on the SMB segment, today.

SOM

\$120M

Our realistic target by year three of scaling.



Why now: capability

Model quality roughly doubles every quarter. A product built on top of models gets better without touching the product.



Why now: cost

API prices fall 40–60% a year. Serving a customer gets cheaper while the price they pay stays the same.




Why now: labour

The talent gap keeps widening. The alternative to us — hiring a human — keeps getting more expensive and harder.

AI AGENT ORCHESTRATION MARKET

\$13.6B (2026) → \$82B (2035) · CAGR 22%



AI agents are crossing from experiments into **everyday infrastructure for business.**

A digital employee, not another tool.

EmAGI gives a business an autonomous AI teammate that lives inside the messenger it already uses.



Persistent memory

Months of context, decisions and preferences. No amnesia between sessions.



Its own server and hands

Own machine, own browser, applications, APIs, CRM, admin panels, ad accounts, documents. It executes, it doesn't advise.



Model-agnostic

Picks the best model per task through MCP — Claude, GPT, Gemini, Grok. No vendor lock-in for us or the customer.



Autonomous

Works on a schedule and on its own initiative, dozens of conversations in parallel, 24/7.



Self-improving

Every night it reviews its own day and writes itself new skills. It gets better at your business specifically.



Quantified value

From \$350 a month against \$4,700 for a person — and the agent doesn't quit, doesn't get sick and doesn't forget.

One sentence: **everything a person can do at a computer, done by an agent that never leaves.**

How it actually works.

Messenger → **Agent** (own machine, own memory) → model chosen per task → **tools and systems** → result back in the chat → **nightly reflection**



Channels

Telegram, Slack, Discord, Teams, WhatsApp, Nextcloud Talk, a widget on the customer's own site. No new interface to learn.



Isolation

Every agent runs on its own machine in its own sandbox with its own memory. One customer's agent cannot reach another's data.



Memory layer

Hybrid keyword and vector recall over the agent's whole history, plus per-chat rules and durable facts about people.



Skills

A shared library plus skills the agent writes for itself. Our longest-running agent has 70 of its own.



Telegram



Slack



Discord



WhatsApp



Teams



Web widget



Nextcloud Talk



Why this is defensible

The hard part is not calling a model — it is keeping an agent alive, remembering, isolated and accountable for months: delivery guarantees, memory that survives compaction, per-customer sandboxes, self-written skills and watchdogs. That is 18 months of production engineering, not a prompt.

Subscription per agent. Simple and high-margin.

STANDARD

\$350 /mo

One agent, memory, core tools and integrations.

Most popular

ASSISTED

\$500 /mo

Hands-on setup, custom skills, priority support.

VIP

\$1,000 /mo

Multiple agents, own server option, full customisation.

\$7.2K

AVERAGE CONTRACT VALUE /
YEAR

90%+

GROSS MARGIN

<5%

API COST OF REVENUE

7–21 days

SALES CYCLE



Gross margin above 90%. Models are billed to the customer's own subscription; our cost to serve one agent is under 5% of its revenue — infrastructure only. Every new customer arrives at the same margin, so growth does not eat the unit economics.

Trial week → real work delivered → paid subscription. **LTV above \$14K at our current retention; break-even at two to three customers.**

In production, paid for, growing without marketing.

Figures taken from our own billing and infrastructure on 29 August 2026.

45

AGENTS LIVE 24/7

19

BUSINESSES

10+

PAYING CUSTOMERS

\$7K+

MONTHLY REVENUE

\$0 spent on advertising

Every customer arrived by word of mouth: people watched our agents work in public chats and asked for one of their own.

A real result, measured

A construction customer sent an agent a hundred site photos and asked for a cost estimate. It was ready in two hours — work their people needed two weeks for.

Depth, not a pilot

Longest-running agent: 480+ days live, 1M+ memory records, 70 self-written skills.

Mars

AI partner for startups (co-CEO):
marketing, financial models,
landing pages, pi

- I see where a landing loses leads and where SEO rankings stand: conversion,



Bob

Building-envelope engineer:
facades, glazing, waterproofing ·
Cyprus

- I track Cyprus dealer prices and base freshness: where a figure has gone stale, I

Asya

Personal research assistant for
home, renovation and purchases
abroad

- I spot a clause in your contract or estimate that tilts against you: an advance with no

Land in the messenger, expand inside the account.



Segments

Owner-run teams, mid-market companies, agencies and niche experts in construction, health, DeFi and e-commerce.



Channels

Where they already talk: Telegram, Slack, X, Discord communities. Our agents are visible doing real work — that is the demo and the ad in one.



Motion

Discover → one-week trial on real tasks → measurable result → subscription. Then a second agent for another function.



Expansion

Q4 2026 US and EU · Q1 2027 enterprise pilots on Teams and Discord · 2027 self-serve platform and an agent marketplace.



What we are buying with this round

Today growth is organic and therefore slow and lumpy. The round converts a proven word-of-mouth motion into a repeatable one: outbound in the US and EU, partner integrators, and a self-serve funnel that does not need us in the loop.

Everyone else built a narrow tool.

	DEVIN	SIERRA	11X.AI	LINDY	CHATGPT	EMAGI
Scope	Code	Support	Sales	Workflows	Assistant	Any business task
Lives in	Web / IDE	Web / CRM	Email / CRM	Integrations	Web / app	Customer's messenger
Memory	No	No	No	No	Limited	Yes
Autonomy	No	No	No	No	No	Yes
Self-improvement	No	No	No	No	No	Yes
Model choice	Locked	Locked	Locked	Locked	Locked	Any, per task

We differentiate on **memory, autonomy and model independence** — and on price: a full-time teammate for the cost of a few hours of a contractor.

The moat is the runtime, not the model.



Core innovation

A production runtime that keeps an autonomous agent alive for months: guaranteed event delivery, hybrid memory that survives context compaction, one isolated machine per customer, a skill system the agent extends itself, nightly reflection, and watchdogs that prove the agent is actually working.



Protection

Trade secrets and know-how; no patents filed. The defensible asset is the accumulated operational knowledge — every failure mode we have already fixed in production — plus each customer's agent memory, which cannot be exported to a competitor.



Barriers to entry

A demo takes a weekend. An agent a business trusts with real work takes 18 months of hardening. Switching cost grows every week: the agent knows the customer's business better than a new hire would.



Scalability

Horizontal by design: one machine per customer, new machines provisioned automatically. Cost to serve stays under 5% of revenue as we grow, and falls as model prices fall.

Three years, built bottom-up from today's numbers.

	TODAY	2027	2028	2029
Paying customers	10+	100	300	700
ARR run-rate	\$84K	\$720K	\$2.2M	\$5.0M
Revenue in year	—	\$350K	\$1.3M	\$3.4M
Gross margin	90%+	90%	91%	92%
EBITDA	—	-\$180K	+\$250K	+\$1.4M
Cash burn / mo	<\$5K	\$18K	break-even	cash generative

Assumptions — the part that matters

- Average revenue per customer \$600/month, mix of the three plans, unchanged over the period.
- New customers: 8/month in 2027, 18/month in 2028, 35/month in 2029 — driven by paid GTM, not by word of mouth as today.
- Monthly churn 3%, flat. Today it is lower, we do not assume it stays there.
- Cost to serve under 5% of revenue; infrastructure scales linearly with customers.
- Headcount stays small by design: our own agents do engineering, billing and support, so opex is infrastructure and GTM, not salaries.

ARR RUN-RATE, BOTTOM-UP

today \$84K



\$84K today → \$5.0M in 2029

\$350K funds 24 months at this burn and takes us **from organic growth to a repeatable sales motion.**

One founder. The rest of the company is the product.

Mykola Kolomyts — Founder & CEO

15+ years building and scaling technology companies. Four years of attempts at an autonomous agent before this one worked. Sole human on the team: sets direction, owns product and customers.



Key hires this round unlocks

A GTM lead for the US and EU, and a second platform engineer. Nothing else — the point of the product is that a company does not need a large team.



The rest of the team are EmAGI agents

Aurora — platform engineering: infrastructure, releases, billing, this presentation and every number in it.

Nova — onboarding: builds and ships new agents to customers.

Leo — representation in founder and venture communities.

They are our own agents, not people. A company of one that ships like a team is the strongest possible proof of what we sell.

Next 12–24 months.

Q4 2026



US & EU entry

Product: self-serve onboarding without us in the loop.

Commercial: first 50 paying customers, outbound in English.

Q1–Q2 2027



Enterprise-ready

Product: Teams and Discord, SSO, audit trail, on-premise option.

Commercial: first enterprise pilots and reference cases.

Q3–Q4 2027



Marketplace

Product: no-code agent templates, marketplace of trained specialists.

Commercial: 300 paying customers, partner integrators.

ONGOING



Data sovereignty

Product: local models as a filtering layer so sensitive data never leaves the customer. The blocker for regulated industries, and it is solvable.

Every milestone is commercial as well as technical: [we ship what customers already ask for.](#)

Raising \$350K pre-seed.

40% — Engineering

Self-serve onboarding, enterprise features, memory and autonomy. One senior hire plus infrastructure work.

40% — Go-to-market

US and EU: outbound, partnerships, first customer-success capacity. This is what turns organic growth into a repeatable motion.

20% — Infrastructure

Machines for new customers, security, reliability and monitoring as the fleet grows.

24 mo

RUNWAY AT PLANNED BURN

300

PAYING CUSTOMERS UNLOCKED

\$2.2M

ARR AT END OF THE ROUND

Break-even

REACHED INSIDE THE ROUND

40% Engineering

40% Go-to-market

20% Infrastructure

We are not raising to find product-market fit. **We are raising to serve demand we already have.**

WHY WE WIN

We are the last generation of managers who will hire people to work at a computer.



Working product

45 agents in production, paying customers, 480+ days of uptime.



Right category

Memory and autonomy — the two things every competitor lacks.



Living proof

A company run by its own agents. We are our own reference customer.

Mykola Kolomyts · founder@emagi.ai · **emagi.ai**

