



Pre-seed • October 2026

Custom software, delivered and maintained.

Brimos is an AI-native service company. Agents build, test and deploy; engineers supervise releases and support. We sell the outcome, not seats, hours or code.

Working MVP

brief to deployed app, run end to end several times

Live demo

debugged with test clients, ready to show

~30 min

from brief to a live app in a clean demo run

A small business needs one workflow built, deployed and kept running.

The gap

Prototype tools generate code.
Freelance developers build to a brief.
Between them sits a gap: someone has to deploy, check and maintain the result, and a non-technical owner usually cannot.

Buyer hypothesis

Non-technical founders and small businesses that need a bot, a web service or an internal tool. First market to test: Russian-speaking founders in Serbia, Cyprus and the EU.

What we are testing

Will these buyers pay a fixed price for one workflow delivered and maintained?

First use case: reminders, status records and outbound API integrations.

The customer agrees the brief. Agents build and ship. Engineers take what agents cannot finish.

01 · Agree

The customer confirms a versioned brief. Requirements are mapped before work starts.

Customer

02 · Build

An architect agent splits the brief into tasks. Coding agents work in isolated containers; every change goes through a pull request and CI.

Agents

03 · Ship and check

Automated deploy, a health check of the service, then a QA pass on the live app.

Agents

04 · Maintain

Monitoring, bounded retries, and follow-up changes through the same brief flow.

Agents and an engineer

When an agent cannot finish, the request is parked and an engineer picks it up.

Automated checks confirm the app is up. Acceptance against the brief is signed off by the customer.

Where we are

The MVP works end to end. We can show it live.

Pre-revenue. Paid demand is the next thing to test, not the product.

Built

The full line from a confirmed brief to a deployed, health-checked app, with monitoring and follow-up changes on the same flow.

Tested

Run end to end several times. Debugged on real requests from test clients.

~30 min from brief to a live app in a clean demo run.

Open for review

Live demo on request. The code behind the product claims is linked in Appendix C.

Why now

Investors are asking for AI companies that sell the service, not the tool.

Y Combinator asked for this

"AI-native companies that don't sell software—they sell the service. Instead of giving you a tool, they just do the work."

YC Requests for Startups, Summer 2026 M9

Services are the larger budget

"For every dollar spent on software, six are spent on services."

Julien Bek, Sequoia, as quoted by TEN13 M3

Agents got good enough

Coding agents now write and fix code inside guard-rails. As they take on more of the build, the hard part moves to hosting, secrets, monitoring and support. We build that layer now.

Market

Buyers already pay for automation work. We size our entry market from pilots.

Entry market, bottom-up

reachable buyers in the first market × annual build and maintenance spend

Both inputs come from discovery calls and pilot pricing, not from top-down reports.

>90% YoY

Growth in Upwork's AI Integration and Automation gross services volume, Q4 2025. Paid automation demand is rising. M2

\$53.5B

Global custom software development, 2026 estimate; 18.8% forecast CAGR for 2026–2033. Category context, not our market. M1

Market signal

The largest app builder is moving toward running products, not only generating them.

"Lovable does not output code. The output is a product, and increasingly so, a business. We do a lot of things around hosting, deployment, and scaling apps."

Fabian Hedin, co-founder of Lovable, September 2026 M10

What it confirms

A company valued at \$13.3B, with more than \$600M in annualized revenue, reached the same conclusion we did: buyers want a working product, not code.

What it leaves open

We will not compete with Lovable head-on. A self-serve builder leaves the user in charge of what ships. Taking responsibility for production for every customer does not scale on that model, so the delivered-and-maintained segment stays open.

Reusable modules alone are not new. We sell responsibility for the result in production.

Modules as contracts

The agent assembles a product from modules defined by a spec, tests and interfaces that can be checked, not from templates.

We own production

Hosting, deploy, encryption, certificates, monitoring and fixes are our job, not the customer's.

People on escalation

Engineers step in when an agent escalates. They do not write every line.

Reuse protects the margin as volume grows. The product we sell is the running app.

Precedents: Crowdbotics built apps from reusable modules with human developers and refocused on legacy-code analysis in 2025 M8. Outsourcing firms sell their own accelerators.

Alternatives

Each option covers part of the job. The gap is who runs the result after launch.

Option	What the buyer gets	Price basis	Runs it after launch
Manual process	Staff time on repeated follow-ups	Not measured yet	The team
Freelance developer	Code built to a brief	Hourly or per-project quote	Agreed case by case
n8n M4	Hosted workflow platform	Starter €20/month, annual billing	The customer
Lovable and partners M5–6	Full-stack hosting; partners build for clients	Platform credits and a partner quote	Customer or partner
Flatlogic services M7	Managed custom development	Monthly team capacity, quoted	Provider team
Brimos	One workflow delivered and maintained	Fixed price and a monthly fee, to test	Our engineers

To prove in pilots: one team owns build, deploy and maintenance at a fixed price. We will report time to launch and support hours per pilot.

Business model

A fixed price for delivery and a monthly fee for maintenance.

To be validated in paid pilots. No revenue is booked yet.

Typical pilot

One workflow. Acceptance criteria come from the confirmed brief.
Fixed price, paid on acceptance.

Proposed band for projects: \$5–30K and 1–4 weeks, to be tested in pilots.

Maintenance

Hosting, monitoring, bug fixes and scoped small changes for a monthly fee.

New features

Quoted separately and built through the same brief flow.

Pilot data stays non-critical and replaceable until backup and restore are tested (gate 0, slide 14).

What one delivery costs, and what reuse has to save.

Cost of one delivery

```
price - (engineer hours × loaded rate + all model  
calls + infrastructure)
```

Every model call and engineering attempt is written to a cost ledger. The pipeline is model-agnostic, so cost is measured on the models we run for clients.

Reuse, and how we will measure it

Reuse base today: two app templates, shared runtime contracts, one catalog package (reminders). Engineers turn generated features into tested modules; agents then assemble new apps from them.

Each pilot records engineer hours and full cost, first delivery against a comparable repeat.

Capacity: 2–3 concurrent projects per engineer is a founder estimate, not a measurement. 4–6 is the target.

First commercial milestone: 3–5 paid pilots, starting from warm channels.

01 · Find the buyer

Warm introductions from founder communities and B2B networks in the first market.

10–20 discovery calls before any outbound.

02 · Sell a pilot

One workflow, acceptance criteria and a defined support term, at a fixed price.

03 · Repeat the scope

A second comparable delivery on the same foundation. Track engineer hours, full cost and support work.

Team

The founder built the platform. Engineering and operations leads join him.

Vladislav Meshkorudnyj

Founder · CEO and CTO

Built the delivery platform alone since December 2025, while employed full-time.

About 7 years of Python backend and AI engineering: Yandex, EPAM, fintech.

github.com/vladmesh ·
linkedin.com/in/vladmesh

Oleg Sergienko

Engineering lead (joining)

Fullstack engineer and engineering manager.

Three-time founder, 300+ B2B clients.

Yulia Kim

Founding operations

7+ years of program and project operations in education, communities and AI Safety.

Led operations for a 150+ participant AI Safety course: 21 groups, 17 mentors, 18 projects.

Hiring by milestone

Backend, ML and frontend engineers join as gates pass. They review agent output instead of writing every line, so one engineer can carry several projects.

Validation plan, 12 months

Four gates, in order. Each one has to pass before the next.

Gate 0 • Readiness

Before any valuable customer data: tested backup and restore, access revocation, TLS on product endpoints.

Gate 1 • Will buyers pay?

3–5 paid pilots, each accepted against its brief.

Gate 2 • Does reuse pay?

A second comparable delivery needs fewer engineer hours at a lower full cost.

Gate 3 • Can we support it?

Pilots renew paid maintenance; support hours per client are measured.

The ask

Pre-seed for 12 months of milestones.

What the round funds

- 01 Readiness gate passed
- 02 3–5 paid pilots accepted
- 03 Repeat delivery measured
- 04 Maintenance offer tested

Founder full-time, engineers joining by milestone, compute.

Milestone for seed

\$200K ARR or a credible enterprise pilot.

Tracked weekly

Projects per engineer; share of each project built from reusable modules.

hello@brimos.dev linkedin.com/in/vladmesh brimos.dev

What a newly generated product can do today, and what comes before paid pilot data.

Supported

- Telegram commands, buttons and menus
- Records kept in PostgreSQL across updates
- Scheduled actions, adjustable settings
- Outbound integrations with customer keys
- Access for named users; ownership transfer

Next in scope

- Web UI, custom domain, Telegram Mini App
- Inbound webhooks, payments, OAuth sign-in
- Built-in email or file storage

Gate 0 work

- Backup and restore of product data, tested
- Owner controls to revoke access and open a bot to the public
- TLS on every product endpoint

Market, category and alternatives.

M1	Grand View Research, custom software development market	2026 estimate; broad category, not our entry market
M2	Upwork, Q4 and FY 2025 results	Demand signal for automation work
M3	TEN13, AI-native services thesis	Quotes Julien Bek (Sequoia) on services spend
M4	n8n, plans and pricing	Starter plan, annual billing; platform price only
M5	Lovable Cloud documentation	Hosted app, database, authentication, storage
M6	Lovable partner program	Partners building and delivering for clients
M7	Flatlogic, custom development services	Managed delivery, monthly capacity, long-term support
M8	Crowdbotics / CoreStory	Rebranded in 2025; now focused on legacy-code analysis
M9	YC Summer 2026 RFS, AI-Native Service Companies	Quoted text of the RFS section
M10	TechCrunch, Lovable crosses \$600M ARR	Hedin quote and revenue; valuation: TechCrunch, 12 Aug 2026

The code behind the product claims is open for technical review.

R1	Product capability map	Scope and limits of newly generated products
R2	Brief and delivery pipeline	Confirmed requirements, escalation, release
R3	Post-deploy smoke checks	HTTP health, bot identity, container state
R4	Product Kit README	App templates, package runtime, composition limits
R5	Package catalog	The reminders package and its versions
R6	Engineering and QA ledger model	Cost attribution; unknown cost is stored as unknown
R7	Testing documentation	Test layers and the limits of controlled evidence
R8	Orchestrator README	Architecture and service inventory

