● 3 CLIENTS IN PRODUCTION

Your product's own assistant.
One that acts.

We give complex B2B software an AI assistant that does the work for the user: through the product's API, in any channel, from the web app to Claude and ChatGPT. So the product stays relevant as users move from apps to agents.

$8,000 / MONTH REVENUE	**3** PAYING CLIENTS, ALL IN PRODUCTION	**1** QUARTER FROM ZERO TO REVENUE	**1** FOUNDER, FULL CYCLE

Users are moving from apps to agents.

The interface is changing hands. People increasingly tell an agent what they want instead of clicking through a product. Analysts expect 40% of enterprise applications to embed task-specific AI agents by the end of 2026, up from under 5% in 2025, and a third of user experiences to move from native apps to agentic front ends by 2028.[1]

Where users are

Claude, ChatGPT, voice, messengers. The product is one tab among many, and increasingly not the tab that's open.

What they expect

Say it once, get it done. Not "find the setting," not "read the doc."

What it means for SaaS

Every product now faces a choice: become something an agent can operate, or become the thing users route around.

CONSEQUENCE The companies that win this shift won't be the ones with the most features. They'll be the ones an agent can drive.

Powerful software is hard to use. That cost is now a risk.

The more a B2B product can do, the longer it takes to reach value. Teams spend months in onboarding. Customer-success people configure the product by hand for every new account. Power features go unused because every capability needs another screen, and the frontend queue never ends.

01

Months to value

"The initial setup took around six months for us."

PUBLIC REVIEW OF A LEGAL CRM

02

Humans as the interface

"A real team runs your switch. Not a help article."

ONBOARDING PAGE OF A PROPERTY-MANAGEMENT PLATFORM; 45 DAYS AVERAGE GO-LIVE

03

Features nobody reaches

"Our team of 8 spent over 200 hours learning the tool and we are still learning."

PUBLIC REVIEW OF A PROFESSIONAL-SERVICES PLATFORM

CONSEQUENCE Value is locked behind the interface. And an agent-first user won't wait for the interface.

An assistant that acts.
Inside the product, through its API.

Not a chatbot that answers from docs. A layer of control: the user says what they want, the assistant does it through the product's own API, shows what it did, and asks when unsure.

USER	"Log two hours to Acme and tell me what's still unbilled on their retainer."
ASSISTANT	picks the right object (retainer, not ticket), creates the activity, returns the unbilled total, with a link to what it changed.

Does, not describes

Creates, configures, changes, checks. Every action logged and reversible.

Honest by design

Refuses when it doesn't know; hands off to a human with full context.

Makes the hard parts easy

Bulk operations, customization, automations, all through conversation. No new UI needed.

WHAT CHANGES Time-to-value in days, not months. Power features finally used. CS stops being the interface.

One assistant.
Every place the user already is.

The same assistant, the same actions, the same permissions, wherever the conversation happens. Channel is a checkbox, not a project.

- 01 **In the web app**
- 02 **In the mobile app**
- 03 **By voice (under 1.5 s to first response)**
- 04 **In messengers**
- 05 **In Claude and ChatGPT via MCP, with per-user OAuth**
- 06 **Listed in the connector directories, where users discover products now**

WHY IT MATTERS Actions described once as tools for the assistant are the same tools for MCP. Ship the assistant and the agent channel comes almost free.

HUMAN IN THE LOOP

Ask for an assistant.
Get one deployed.

Every client assistant is assembled on one core by our meta-concierge: describe the product and the flows, and it builds the assistant and deploys it. Ask for a change, and it changes it. No core rewrites, no project from scratch.

01 **Knowledge**
The client's knowledge base with versions. The assistant cites the source, not the model's memory.

02 **Tools**
The product's API, other products via their MCP servers, and scripts generated for a user's own workflow.

03 **Access ceiling**
The assistant never sees more than the human it's talking to. Enforced in code, not in the prompt.

04 **Verification**
Every answer and action passes a fact-and-tone gate before the user sees it.

\+ **Skills**
Reusable flows on top of tools. Each user ends up with their own set: the product, customized around how they work.

CONSEQUENCE A new client assistant is configuration, not a new project. That's why three clients shipped in one quarter with one person.

LIVING

It gets more accurate every week.

Shipping is the start. Every conversation is analyzed; an eval suite runs against real usage; each week the assistant gets a retro and an improvement: a new flow, a better tool, sharper knowledge.

Every conversation

Analyzed for gaps, misfires, missing actions.

Weekly

Retro → fix → redeploy. One new flow a month.

Eval report

What it closed alone, what it handed to a human, what to fix next. As a live page, not a PDF.

CONSEQUENCE Most AI features are launched and left. Ours is maintained like a product, because it is one.

Three paying clients.
All in production.

$8,000 / MONTH	**3** CLIENTS	**0** PILOTS THAT FAILED TO SHIP	**1** PERSON ON THE TEAM

PRODUCT COMPANY

Assistant over the product.

Integration layer plus an assistant that works with the product, not about it; MCP layer for agent channels.

B2B SERVICES

R&D engine.

Internal product with live AI-driven reports instead of static PDFs; team asks questions by voice. AI-readiness audit accepted; moving to company-wide transformation.

200-PERSON COMPANY

CAIO as a service.

Contracted Dopomogai as fractional Chief AI Officer: a living map of every business process and how AI overlays it, then roadmap and implementation. Ongoing retainer.

Built and shipped by one person

Margin isn't eaten by a team.

Three different engagements on one core

Product assistant, internal process assistant, company-wide transformation. A platform with a ladder, not project outsourcing.

Every client came by referral

Marketing hasn't started yet.

WHAT IT PROVES Demand existed before the first dollar was spent finding it, and the ladder from product to company already has a client on every rung. Client names on a call.

Businesses are re-platforming on AI. We're the layer that makes it usable.

Worldwide AI spending reaches $2.59 trillion in 2026, up 47%, and Gartner calls 2026 "the inflection year" when enterprises, not hyperscalers, start to spend.[5] Of that, about $1 trillion is AI software and services, the money companies pay to put AI into how they work.[5]

THE MARKET WE COMPETE IN

AI-ization of business

~$1T/year in 2026 on AI software and services.[5] Every company will re-platform its products and processes around agents; most don't know how. This is where CAIO-as-a-service and process assistants live.

THE SEGMENT WE'RE IN

Agents inside business software

40% of enterprise apps embed task-specific agents by end of 2026, from <5%.[1] ~30,000 SaaS companies worldwide, ~17,000 in the US.[2] MCP, the open standard for agents operating products, passed 10,000 public servers within a year.[3]

WHERE WE START

The wedge

US B2B SaaS, 15–150 people, complex product, public API, founder still decides. Our own scan: of 62 such companies, a third had already shipped an MCP server, and none were listed where users look.[4]

HONEST MATH We don't need a percentage of a trillion. 100 product clients on today's pricing is ~$6M/year recurring before builds, and the process layer behind them is bigger than the product layer. The constraint is our sales capacity, not demand.

Products first.
Then the whole company.

An assistant that operates one product is the wedge. The same core (knowledge, tools, access ceiling, verification, living loop) runs assistants over any process a company has. We're already doing it.

NOW

Product Concierge

Assistant inside a SaaS product + agent channels. What we sell in the US today.

NEXT

Company processes

Assistants over internal operations: R&D, sales, support, onboarding. Already live at one client.

THEN

CAIO as a service

A living map of every business process and how AI overlays it; roadmap and implementation as a retainer. The company doesn't hire a Chief AI Officer. It subscribes to one, with the platform underneath. Already sold once: a 200-person company has contracted Dopomogai as its fractional CAIO to map and re-platform every process. That map is becoming the product.

LOGIC The product assistant is the cheapest way to prove we can be trusted with a business process. Everything else is the same core, more processes, and two of our clients are already there.

Build.
Then keep it alive.

SERVICE: WE DO IT

Build

$12k/month × 3 months. Month 1: first real flow with actions, live for their team. Month 2: 3–5 flows, branded widget, MCP with per-user OAuth, submitted to directories. Month 3: public, voice if needed, eval report.

Living

$4–6k/month. Weekly analysis, evals, improvements, one new flow a month.

Platform

From $300/month, bring-your-own-keys.

Then

AI roadmap → implementation retainer (CAIO as a service).

PRODUCT: THEY DO IT

Console

Self-serve: build and deploy your own assistant on the same core. Subscription + usage. In development; first channel: e-commerce (Shopify).

What one client is worth

Build ~$36k, then Living ~$5k/month: ~$80k in the first year, ~$60k/year recurring after. Process and CAIO work on top multiplies it.

Why now

The agent channel is being standardized right now; whoever ships the assistant first owns the customer's interface.

CURRENT PRICING One Build client ≈ $80k first-year value, ~$60k/year after; margin protected by the meta-concierge doing the assembly.

One founder.
Full cycle, from code to client.

10 years
In engineering and product management. Knows both sides of the table: how it's built and how it's sold.

Writes it
Architecture, backend, frontend, agents, deploy. Idea to production without intermediaries.

Ships it
Leads every client from first call to launch. Knows exactly where the product breaks against reality.

RISK, HONESTLY A solo founder is a bottleneck in sales, not in product. That's what the next hire and this competition are for.

Three things.
In this order.

$8,000	3	1	solo
MRR	PAYING CLIENTS	QUARTER	FOUNDER

01 **First US clients**

Introductions to founders of complex B2B SaaS (15–150 people, public API). We build a working preview on their sandbox before the first call.

02 **Mentors and network**

People who've sold to US SaaS and know where a European approach breaks.

03 **Capital for the first hire**

Grant or prize funds go to the first engineer, so the founder stops being the only sales channel. Not raising a priced round today; open to the conversation.

Ask this page anything.

The assistant in the corner is built on the same core we sell. Ask it about pricing, the cases, the architecture, or ask it to explain the shift to your CFO. It cites its sources.

Andrew
Founder, Dopomogai
solovei.andrew@gmail.com LinkedIn dopomogai.com

Sources.

Numbered as cited on the page. Where a source is secondary, it says so: those numbers are being verified against the primary before we quote them to investors.

1 Gartner press release, 26 Aug 2025: "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026".

2 Statista via Ascendix ("How Many SaaS Companies Are There"), early-2025 count. Secondary source, being verified before we quote it to investors.

3 Anthropic MCP ecosystem update, Dec 2025 ("more than 10,000 active public servers"), cited via DigitalApplied. Secondary source, primary being verified.

4 Dopomogai desk research, Sept 2026 (62 companies; 20 with official MCP; 0 in the Claude directory).

5 Gartner press release, 19 May 2026: "Gartner Forecasts Worldwide AI Spending to Grow 47% in 2026". $2.59T total; AI services $585.5B; AI software $453.2B.

6 IDC press release, 2025: "Agentic AI to Dominate IT Budget Expansion Over Next Five Years". Worldwide AI spending $1.3T by 2029, agentic AI the main driver, ~10× increase in agents used by enterprises. Read via summary; secondary source.