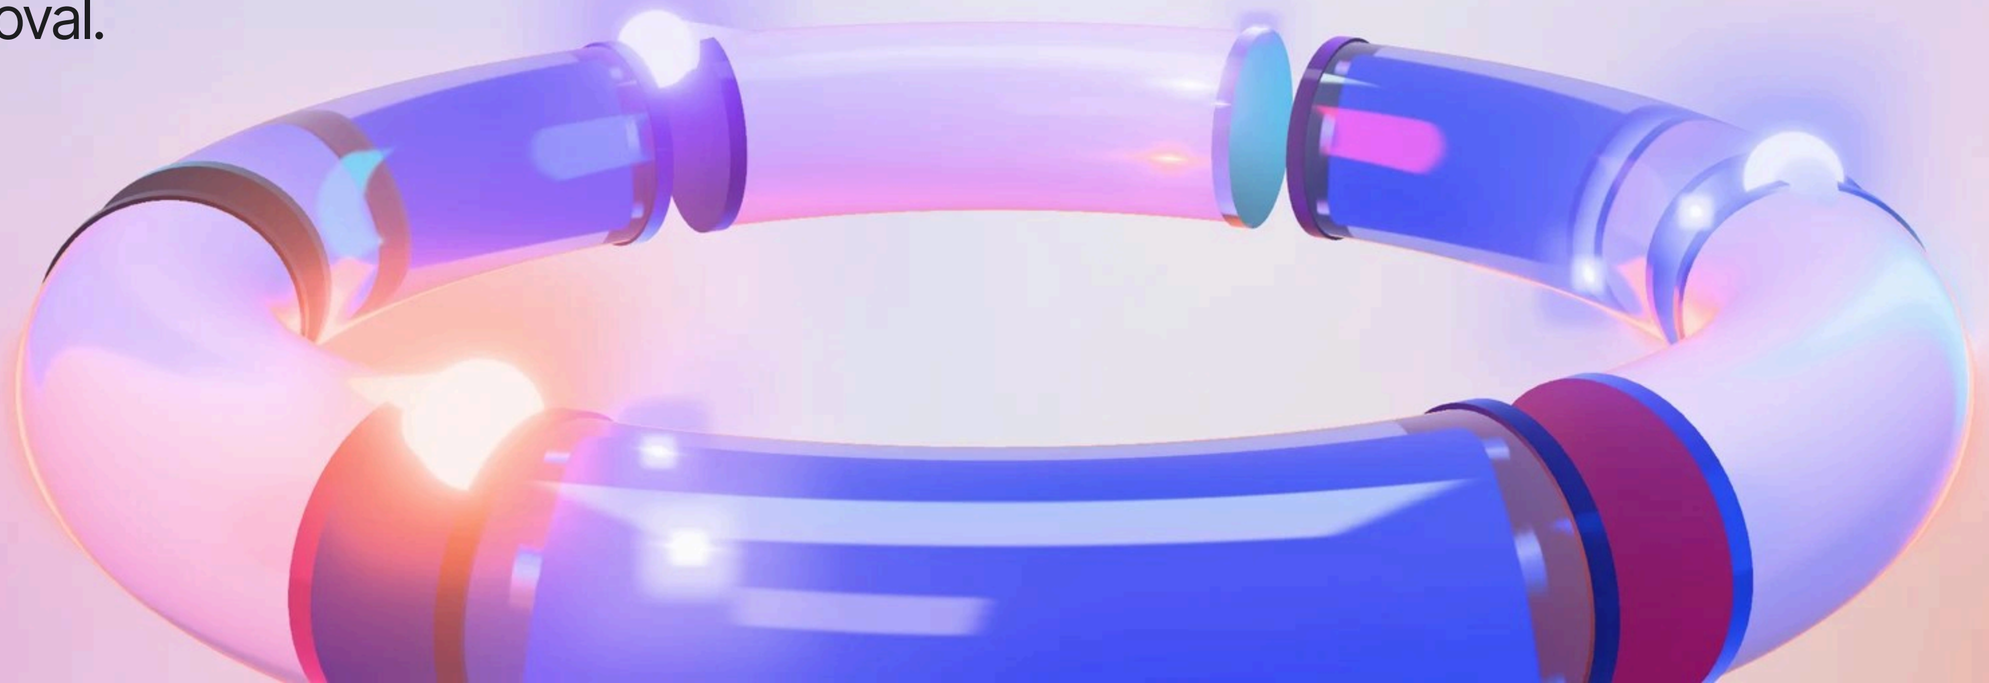
INVESTOR OVERVIEW · NOVEMBER 2026

Vitalii Cherepanov

Status as of 4 November 2026

LOOPmKit

A builder of AI agents that you run on your own server. Agents work in a loop, inside a budget, and risky actions wait for your approval.



THE PROBLEM

Agents are easy to start and hard to trust.

Founder	“Every task lands on me.” “The AI bill is different every month.”
Operations lead	“A failed automation runs twice, or fails in silence. What did the agent actually do?”
Marketing lead	“Nothing goes out under our name until a person says yes.”
Engineering lead	“I do not want to be tied to one model vendor.”
Privacy-sensitive team	“Our data stays on our own servers.”

The pains the product was designed against, from its own design notes, in first person for brevity. These are not interview quotes.

THE IDEA

One job per agent. A record wakes the next.

LOOPmKit lets you draw small AI agents on a canvas. Each does one job and writes a checked record. Other agents wake on that record, so the work goes around in a loop.

01

Trigger

A schedule, a webhook, a new record, another agent finishing, or a click.

02

Agent

One job. Typed inputs, a model, tools and a memory of its own.

03

Record

The result is checked against a schema and saved, or retried.

04

Gate

If, filter, delay, join, limit: eight kinds, drawn on the canvas.

05

Action

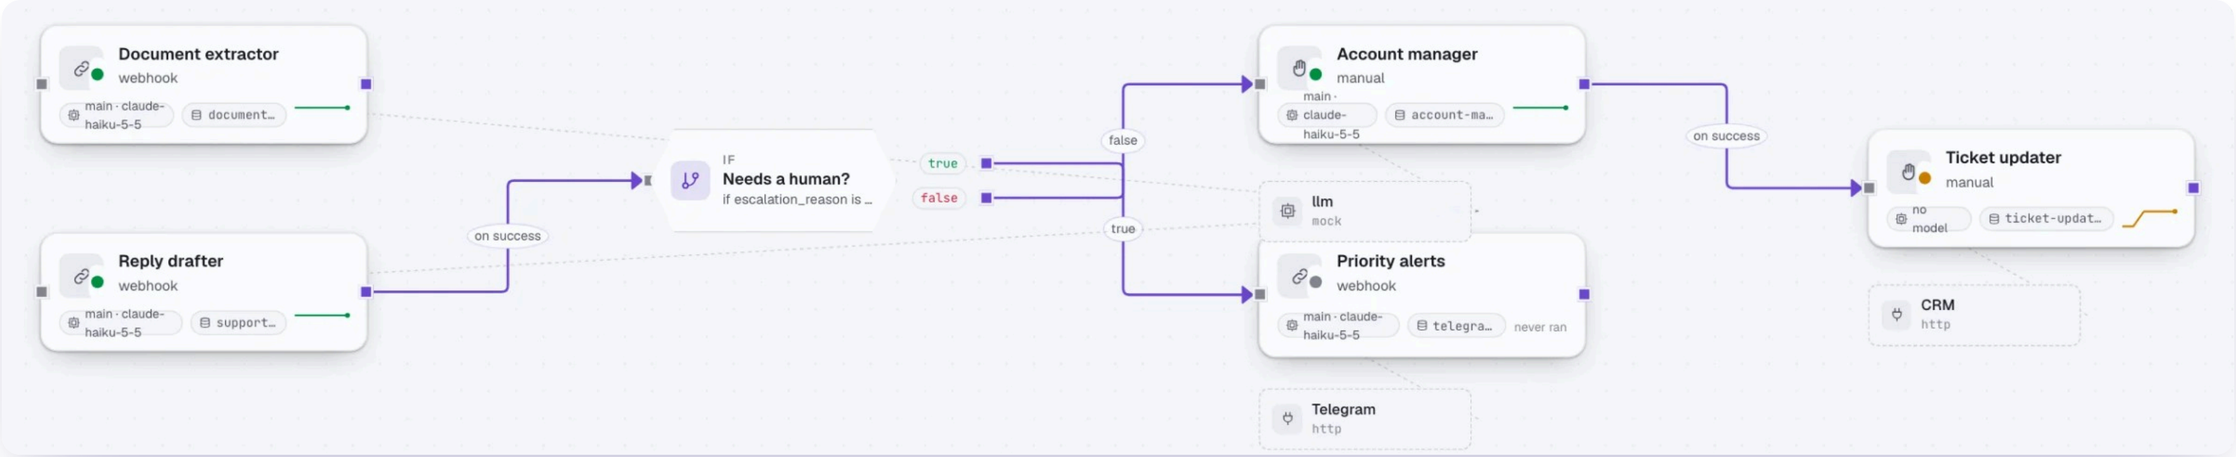
Read, draft or act. A draft waits for a person.



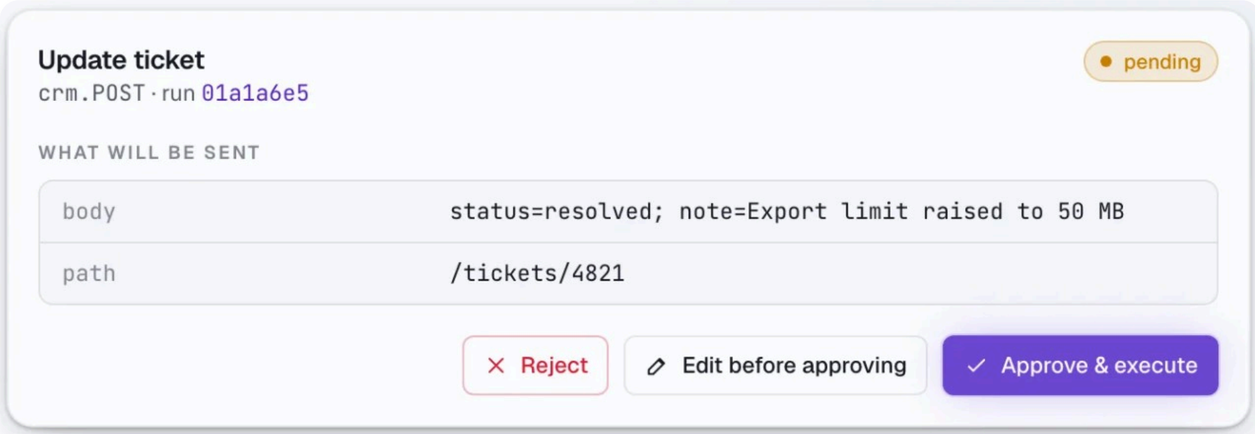
The usual hand-off between agents is a record, not a call. Every step is logged.

THE PRODUCT

The canvas, and the approval.



Canvas: agents, a gate and links between them, wired by outcome.



A drafted action waits in a queue. A person rejects it, edits what will be sent, or approves it.

Trust ladder: read runs, draft waits, act runs now. The lower of the action and the connector scope applies.

Screens come from a test project, not from a customer.

CONTROL AND TOKEN ECONOMY

Every run has a price limit and a gate.

Budgets

A cap per run, per agent per day and per project per day. Over the cap, the run is skipped.

Approvals

Risky actions are drafts first. A person approves, edits or rejects before anything is sent.

Dedupe

An input already seen, by hash, is skipped before any model call.

Caches

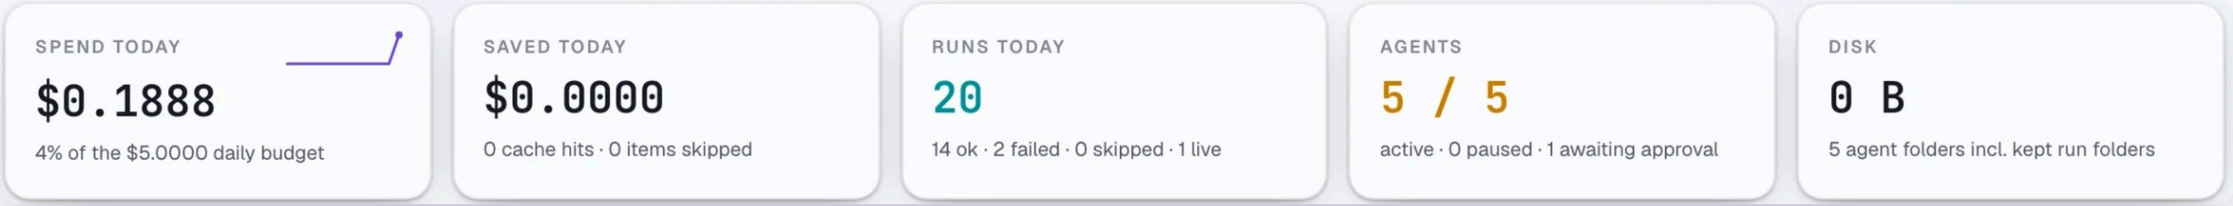
An exact-match response cache, and provider prompt caching of the stable part of a prompt.

Routing

The cheapest enabled model that has the tools, vision, context and quality the task needs. It escalates on failure.

Auto-pause

An agent that fails three runs in a row pauses itself until someone looks.



Overview screen of a test project: spend against the daily budget, cache hits, items skipped.

MODULES

Modules, and how far each has gone.

● has run on a live stack ◐ built and tested, not run live

PLATFORM

● Agent runtime and triggers
5 trigger kinds, queue in PostgreSQL, resume from checkpoint

● Budgets, dedupe, caches
Working on the founder's instance

● Questions and approvals to the owner
Question, answer and resume checked live

◐ Model routing
25 provider presets, local models; learned cascade

◐ Memory per project
Private by project, audience levels

◐ Connectors
61 presets; listed tools, not each one proven

MORE MODULES

◐ Quality checks (evals)
10 evaluator kinds, golden datasets, replay

◐ Company wizard
Checked in a browser with a scripted model

● Video production
Script, voice, subtitled render. One 17-second clip made live

◐ Telephony
Inbound AI receptionist built and tested against a simulated carrier. No real call yet. No outbound calls, no SIP, no mobile app

◐ Social and marketing
14 connectors, 7 networks. Content calendar run as a dry run. No real post through the product yet

◐ Developer flow
Isolated jobs for Claude Code and Codex. Never run with a real vendor subscription

ROLES

29 roles, mostly text-only so far.

● has run live ◐ needs code execution or real accounts, not connected yet ○ text work, tested

MANAGEMENT	Secretary, Chief of staff
PRODUCT AND ENGINEERING	Product manager, Project manager, ◐ Tech lead, ◐ Backend Go, ◐ Backend PHP, ◐ Frontend Vue, ◐ DBA, Mobile, QA, DevOps, Designer, Security
GROWTH	Analyst, Marketer, Content writer, ● Video producer, ◐ Social media, Account manager, Support
FINANCE AND LEGAL	Accountant, business, Accountant, personal, Lawyer, Procurement, HR
KNOWLEDGE AND FRONT DESK	● Researcher, Knowledge keeper, ◐ Receptionist

The founder's own build has 30 active roles; one personal-affairs role is not shown. Role prompts say it plainly: they work with text and do not yet run code or touch real accounts.

TWO WAYS TO START

One agent, or a whole company of roles.

8

task templates

Separate agents

Start from one of 8 task templates, or describe the job in plain text. Each agent has its own mission, schedule, memory and budget.

Mail sorting, a morning digest, research with sources, a document to data fields, a customer reply draft.

10

default roles

A company of roles

A wizard creates the team in one step: 10 default roles, up to 32. Each employee is a project with its own budget, memory and skills.

Employees call each other as functions. The org chart is a set of share policies. Everything is journaled.

The gallery holds 23 templates: 8 task scenarios and 15 employee roles.

WHAT IS BUILT

What exists today, counted.

155,675	lines of Go, in 782 files	**61**	connector presets, 14 of them social
154,375	lines of Go tests, in 651 files	**25**	model and voice provider presets
104	SQL migrations	**23**	gallery templates: 8 tasks, 15 roles
8	kinds of logic gate	**10**	kinds of quality evaluator

How counted: by reading the product repository on 4 November 2026. Lines are newline counts of files in the working tree, without dependencies and generated files. Presets, templates and gate kinds are counted from the product's own lists. Nothing was run to produce these counts, and no test results are claimed here.

REAL USE

Running on its founder's own instance.

37,292 scheduled runs since 27 October 2026



34,707
skipped, nothing new

2,417
succeeded

166
failed

Most scheduled work is the system correctly finding nothing new to do. The failed runs are 166 of 37,292, under half of one percent.

projects	**37**
agents, 175 enabled	**197**
records stored	**26,877**
tokens used	**62.7 M**
duplicates skipped before any model call	**52,410**
paid API spend, excluding subscriptions	**$13.50**

Snapshot of 4 November 2026. Includes test projects. There are no customers. The remaining 2 runs were cancelled.

MARKET AND WHY NOW

Adoption is real. So is the failure rate.

THIRD-PARTY FIGURES AND FORECASTS, NOT OUR RESULTS

62%

of organizations are at least experimenting with AI agents; 23% are scaling an agentic system somewhere.

McKinsey, The state of AI in 2025, November 2025; 1,993 respondents

40%+

of agentic AI projects will be canceled by the end of 2027, Gartner predicts: rising costs, unclear value, inadequate risk controls.

Gartner, press release, 25 June 2025 (prediction)

Cost and risk control are named as reasons agent projects fail. Budgets and approvals are what LOOPmKit is built around.

McKinsey: https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai
Gartner: https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027

LANDSCAPE

Others lead on reach. We bet on control.

	LOOPmKit	n8n	Zapier Agents	Dify	Flowise†	CrewAI	LangGraph	Relevance AI	Lindy	Copilot Studio
Self-hosting	part.	yes	?	yes	yes	yes	yes	?	?	?
Visual canvas	yes	yes	?	yes	yes	yes	part.	?	?	yes
Hard budget per run, agent, project	yes	?	part.	?	?	?	?	?	?	part.
Routing to the cheapest suitable model	yes	?	?	?	?	?	?	?	?	?
Read, draft, act trust levels	yes	?	?	?	?	?	?	?	?	?
Approval queue, arguments editable	yes	part.	part.	part.	?	part.	part.	part.	part.	part.
Dedupe and cache savings	yes	?	?	?	?	?	?	?	?	part.
Connectors	part.	yes	yes	?	?	?	?	yes	yes	yes
Telephony	part.	?	?	?	?	?	?	yes	yes	yes
Hosted cloud	no	yes	yes	yes	part.	yes	yes	yes	yes	yes
Certification (SOC 2 or similar)	no	yes	yes	part.	?	?	?	yes	yes	yes
Customers	no	yes	yes	?	?	?	?	yes	part.	yes

Vendors' own public pages, read 4 Nov 2026. ? = not found in the public docs we read, which is not the same as no. LOOPmKit column is from our own audit; yes means built, and some built parts are not switched on in every setup. part. = private repository without a license, telephony never run live, connectors only listed. We claim no superiority. † Flowise: repository archived by its owner on 13 Aug 2026 (GitHub). Sources: SOURCES.md.

HONEST STATUS

Where we are not there yet.

The “AI company” goal is not achieved.



6 of 10 acceptance criteria are open:

- Employee management with a test run
- Checks that connections are authorized
- The full development process
- The support process
- An owner cockpit
- Full test and live acceptance

No customers and no revenue.

One operator. No multi-tenancy, no roles and rights.

Most roles work with text only.

No real phone call and no real post yet.

Git history starts on 30 October 2026. By dated documents, development is about two weeks old.

License undecided. The repository is private.

The goal is the founder’s own checklist for a company of AI roles. The product’s own documents say the goal must not be declared achieved.

ROADMAP PLAN, NO DATES

Built from what is unfinished.

Plan 1

Prove it end to end

- Close the six open criteria
- A first real phone call
- First real posts, through approvals
- One development run on a real vendor subscription

Plan 2

Switch on what is built

- The product’s own MCP server
- OpenTelemetry tracing
- Content guardrails
- Price database and mandatory memory

Plan 3

Make it installable by others

- An employee catalog with versions
- One-command installer, second version
- Move the founder’s own scripts into the product
- Decide the license and the multi-user model

Derived from the open items in the product audit and its backlog. It is a plan, not a promise, and no dates are committed. The license and multi-user decisions are open questions, not backlog items.

FOUNDER

Built by one person.

Vitalii Cherepanov

Software developer. Writing software since 2002. Novi Sad, Serbia.

Builds open-source infrastructure for AI agents. LOOPmKit uses total-agent-memory for the memory of each project.

total-agent-memory

Persistent local memory for AI coding agents over MCP. MIT license.

74 stars 35 releases latest v14.10.0, 29 Oct 2026

a2abridge

An open A2A 1.0 mesh for Claude Code, Codex, Cursor, Cline, Continue and Gemini. MIT license.

16 stars 10 releases latest v4.0.3, 29 Oct 2026

Source: vbcherepanov.com and the public GitHub API, counted on 4 November 2026.

NEXT STEP

Let us talk.

A 20-minute intro call. Bring your hardest questions.

cal.com/vbcherepanov/intro

Amount and terms are to be discussed in person.

This deck does not state them.

