



// A REPOSITORY RADAR SPECIAL REPORT

OSS Vector

Summer 2026 · the state of open-source software 📡

Open source is the fastest mechanism software has for evolving its own stack: millions of projects compete for adoption, and developers and users collectively pick the winners. This is where that selection pointed in the summer of 2026, read across 242 high-signal repos we hand-picked for importance, momentum, and novelty, then tracked across stars, commits and downloads while they moved.

242

REPOS HAND-TRACKED

8.3_M

STARS UNDER WATCH

914_k

FORKS

67_M

PACKAGE INSTALLS / WEEK

14_k

COMMITTS / WEEK SHIPPED

11/100

OF GITHUB'S ALL-TIME
TOP

This is a curated read, not a scrape

Everything here comes from a deliberate, opinionated sample tuned for momentum and novelty: the moving edge of the field, a high-signal shortlist, not a census of GitHub.

THE METHOD

Every two weeks we hand-pick **~7 repositories** worth your attention, then keep tracking them long after, watching each one rise or stall in real time. Judgement, not a trending-page scrape.

And we follow far more than stars: commit velocity, contributors, issue health and downloads, so "popular" and "actually used" can be told apart.

CHOSEN FOR THREE THINGS



Buzz

what's breaking out right now



Real impact

what teams actually adopt



Our read

where we think it's heading

⚖️ DIRECTIONALLY RIGHT, NOT PERFECT

242 repos across 38 sweeps and ~18 months. A curated lens misses things and carries our bias, but it surfaces signal a firehose buries. Read it as a shortlist, not a census. **All figures as of 16 July 2026.**

Four seats at the same table

Read that selection closely and you're reading the field's leading indicator: open source is where AI adoption and the software-market contest now actually play out. The same repos, four ways, depending on where you sit.



Builders

ENGINEERS & FOUNDERS

See what's already commoditized before you rebuild it. The harness, the gateway and the token-cost layer are solved. Build on them, don't re-pour the foundation.

take: know what's a dependency, not a project



Operators

ENG LEADERS & PLATFORM

See what's real enough to deploy (maintained, actively shipped, actually downloaded), and the new budget line nobody had a year ago: what your agents cost to run.

take: adopt on health, not hype



Investors

THE CAPITAL READ

See where stars, usage and capital point the same way, and the structural shift from libraries to end-user products that resets who the buyer is.

take: the app layer is the new thesis



Researchers & the curious

THE FIELD READ

See the shape of open source right now: what it's made of, how fast it moves, who builds it, and how concentrated (and fragile) the frontier really is.

take: the ground truth, quantified

The report in three parts

A quick map before the deep end: **9 readings**, **2 tables** of the field, **3 calls** we'll be graded on. Everything that follows sits in one of three parts.

I

p. 05-14

The readings

Nine things that define the space right now: the deep end, one reading to a page.

it's all AI

agents are the field

memory splits from RAG

libraries → products

the vertical climb

coding agent = platform

measure-and-run layer

stars ≠ the story

frontier is 1-2 people

II

p. 15-16

The field, quantified

The descriptive base under the readings, and the leaderboards, ranking what actually moved four ways.

languages

licenses

who's behind it

leaderboards x3

III

p. 17

The outlook

Where open source and AI go next, and the specific repos our live signals flag before the crowd does.

the through-line

3 calls we'll be graded on

repos to watch

// THE HEADLINE 📅

Nine readings that define the space right now

Every reading on one page. The next pages take the big ones apart, one at a time.



READING 1

It's all AI now

The "state of open source" is really the state of open-source AI. The non-AI repo is the exception.

~9 in 10 touch AI



READING 2

Agents *are* the field

Not a category among many. The single biggest bucket, and everything else increasingly serves it.

92 repos • 38%




READING 3

Memory splits from RAG

Persistent agent memory and code-as-a-graph broke out of the RAG bucket into their own category this year.

a new line item



READING 4

Libraries → Products

The winners ship as applications you run: desktop apps, self-hosted web apps and binaries, not just a library you import.

39% ship a runnable build



READING 5

The climb went vertical

Cold start to six-figure stars now takes months, not years. The all-time leaderboard is being rewritten in real time.

OpenClaw: 100× in 6 months



READING 6

The coding agent is a platform

A whole harness layer now runs, switches and multiplexes coding agents, plus a token-cost economy underneath.

harness + cost layers



READING 7

The measure-and-run layer

As agents hit production, the tools to evaluate them and to run them reliably matured into their own infrastructure.

eval + gateways



READING 8

Stars are not the story

Buzz and usage disagree. Some of the most-downloaded repos are nowhere near the top of the star board.

19M d1 on 26k stars



READING 9

The frontier is 1–2 people

Behind the hype, most of these projects are carried by a single dominant author.

42% bus-factor 1

// DEEP DIVE • READING 1

It's all AI now

The "state of open source" is really the state of open-source AI. Across all 242 repos, nearly every category is a different face of the same build-out.

CATEGORY • REPOS COVERED

 Agents & Orchestration		92 38%
 Coding & Dev Tools		29 12%
 Infra, Data & Ops		27 11%
 RAG & Memory		21 9%
 Voice, Vision, Multimodal		21 9%
 Models & Inference		20 8%
 Other		17 7%
 Search & Research		11 5%
 Security		4 2%

THE ONE-LINE READ

~9 in 10

REPOS TOUCH AI

Strip out a thin tail of pure-infra and security tools and the field is, functionally, **open-source AI**. The non-AI repo is the exception now, and even the plumbing (gateways, memory, inference) exists to serve models and agents.

Read the split on the left as one thing wearing nine hats: nearly every bucket is a different face of the same AI build-out.

// DEEP DIVE • READING 2

Agents are the field

Not one category among many. Agents & Orchestration is the single biggest bucket, and everything else increasingly serves it.

THE BIGGEST BUCKET, BY FAR

92

AGENT REPOS

38%

OF THE FIELD

Agents & Orchestration is bigger than Coding, Infra and RAG *combined*, and the boundary keeps expanding. Memory, gateways, voice and inference increasingly exist to serve agents, not to stand alone.

"Agentic" stopped being a category this year and became the field's organising principle: the thing most other repos are built to feed.

★ SUPERSTAR OF THE FIELD



hermes-agent

"The agent that grows with you." The clearest single embodiment of the reading: a pure agent, and the field's second-fastest mover, from a standing start to 215k stars in months.

215k

STARS

+57k/mo

PEAK VELOCITY

2.0k

CONTRIBUTORS

#2

BY VELOCITY

HERMES VS THE MEDIAN REPO ON OUR LIST

star velocity		51x
commits / wk		127x
contributors		29x

Behind it the bucket splits three ways: orchestration frameworks ([CopilotKit](#), [Mastra](#)), autonomous agents ([agency-agents](#), [ELIZA](#)) and shipped assistants ([OpenClaw](#)).

Memory splits from RAG

For two years "RAG" was the catch-all for anything that fed context to a model. This year it split: **persistent agent memory** and **code-as-a-graph** broke out as their own category, distinct from plain document retrieval.

BREAKING OUT OF THE RAG BUCKET • BY STARS

claude-mem	context across sessions	87k	★
Understand-Anything	code → knowledge graph	74k	★
codegraph	code as a graph	60k	★
agentmemory	persistent agent memory	25k	★
cocoindex	incremental live index	11k	★

WHY IT'S ITS OWN THING NOW

Retrieval answers "what is relevant to this query." Memory answers "what does this agent already know." Once agents began running continuously rather than one-shot, the second question stopped being a RAG setting and became its own product.

Code-as-a-graph is the same move for engineering: instead of chunking files for retrieval, the whole codebase is indexed as a structured graph the agent can traverse. The fastest-rising shape of it, [codegraph](#), added +43k stars a month.

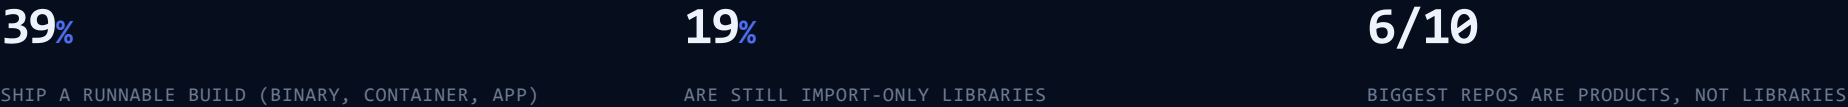
THE SHIFT

Memory is no longer a feature inside RAG. It's a category with its own breakout repos.

// DEEP DIVE • READING 4

→ Libraries → Products

Open source used to mean a **library** you **import**. The breakout repos now are **applications** you **run**: desktop apps, CLIs, self-hosted web apps and native binaries, aimed at end users, not just developers.



PRODUCTS, NOT JUST LIBRARIES • BY STARS

 OpenClaw	AI assistant • any-OS app	383k
 Open WebUI	LLM UI • self-host	145k
 Immich	photos • self-host	108k
 meetily	notetaker • desktop	24k
 orca	agent IDE • desktop	20k
 Reor	notes • desktop	9k

LIBRARY VS APPLICATION

The distinction is concrete: a **library** is code other developers import (npm, PyPI); an **application** is a finished tool an end user runs (desktop app, CLI, self-hosted web app, native binary). Counted by what each repo actually ships: **39%** now ship a runnable build, just **19%** are import-only libraries, and the rest run from source.

That resets the competition: an open-source project is no longer "a package on npm," it's a product going head-to-head with a funded SaaS, on distribution, polish and trust, not just code.

THE SHIFT

Open source went consumer. The unit is the app, not the import.

// DEEP DIVE • READING 5

From cold start to six figures, in months

A GitHub project once took years to reach six-figure stars. Now it takes months. Of the 14 repos we track above 100k, **nine are less than two years old**, and the steepest went roughly **100x in half a year**. Here's the 100k+ club and how fast each climbed.

THE 100K+ CLUB • FROM → NOW • TIME SINCE COVERED

OpenClaw	4k → 383k	6 mo
hermes-agent	27k → 215k	3 mo
OpenCode	33k → 186k	8 mo
Langflow	60k → 152k	14 mo
Firecrawl	23k → 151k	17 mo
Open WebUI	89k → 145k	15 mo
Claude Code	5k → 138k	16 mo
agency-agents	73k → 132k	3 mo
spec-kit	29k → 121k	9 mo
llama.cpp	90k → 120k	8 mo
cc switch	5k → 117k	8 mo
Immich	82k → 108k	8 mo
Gemini CLI	52k → 106k	12 mo
BrowserUse	29k → 105k	17 mo

THE STEEPEST CURVE IN THE FIELD



OpenClaw: **4k → 383k in six months**, a hundredfold rise and the steepest curve we've tracked. Cold-start to viral is now a matter of weeks, not years, and eleven of these are among the 100 most-starred repositories in GitHub's history.

// DEEP DIVE • READING 6

The coding agent became a platform

A whole stack has grown on top of the coding agent: four layers, none of which existed in this form a year ago. The agent itself is now the least interesting part.

THE HARNESS LAYER

Run, switch, multiplex the agents

Once the coding agent is a commodity, value moves to orchestrating several at once, on the accounts you already pay for.

cc switch 117k

Ruflo

emdash

herdr

Superset

THE TOKEN-COST ECONOMY

Tools that exist to cut the bill

A cluster that did not exist a year ago, built purely to shrink what agents cost to run. Per-token prices held flat, but agents burn so many tokens per task that the bill became the product.

headroom +38k/mo

RTK

Context Mode

codegraph

CodeBurn

THE MCP LAYER

How agents reach tools and data

The connective tissue has standardised, and it is now the most-installed corner of the whole field: **FastMCP alone pulls 19M downloads a week**, on a fraction of the stars.

FastMCP 19M dl

Playwright MCP

chrome-devtools-mcp

SKILLS & MARKETPLACES

What the agents can actually do

The top of the stack: packaged capabilities and the marketplaces distributing them: the app store forming around the agent.

Agents marketplace 38k

taste-skill

agents-cli

// DEEP DIVE · READING 7

The measure-and-run layer grew up

Once agents shipped to production, two support layers matured fast: the tools to **measure** them (evaluation, tracing) and the tools to **run** them reliably (gateways, serving). Both are now infrastructure, not afterthoughts.

MEASURE · EVAL & OBSERVABILITY

MLflow	eval + experiment tracking	10M dl/wk
Deepeval	LLM eval framework	1.2M dl/wk
TensorZero	LLMOps platform	12k 

RUN · GATEWAYS & SERVING

LiteLLM	one API → 100+ models	54k 
OmniRoute	free AI gateway	18k 
FlashInfer	LLM serving kernels	1.1M dl/wk

The tell that a field has grown up: "did it work?" and "can it scale?" become products of their own. Evaluation is now a downloads heavyweight (MLflow alone pulls **10M installs a week**), and the gateway/serving layer is the plumbing every production agent quietly runs on.

// DEEP DIVE · READING 8

★ ≠ 📦 Stars are the headline, not the story

A star is a bookmark: it measures attention, not use. Rank the field by what people actually install or download instead, and the order barely holds.

📦 LIBRARIES · WEEKLY INSTALLS (NPM / PYPI)

LiteLLM	170M / wk	d1#1 · ★#14
FastMCP	18M / wk	d1#2 · ★#31
BrowserUse	11M / wk	d1#3 · ★#6
MLflow	11M / wk	d1#4 · ★#29
OpenAI Agents	8.1M / wk	d1#5 · ★#27
Playwright MCP	5.9M / wk	d1#6 · ★#22

💻 PRODUCTS · BINARY DOWNLOADS (GITHUB, TOTAL)

Voicebox	1.2M	d1#1 · ★#24
OpenCode	1.1M	d1#2 · ★#3
cc switch	650k	d1#3 · ★#6
onlook	435k	d1#4 · ★#39
RTK	416k	d1#5 · ★#14
KittenTTS	230k	d1#6 · ★#56

Two usage signals, one story: weekly installs rank the libraries, binary downloads rank the products, and neither tracks stars: LiteLLM takes 170M installs a week at #14 on stars (automated CI builds inflate the raw count, but the direction is clear), Voicebox leads product downloads at #24. Usage numbers exist only where they're public (weekly installs for npm/PyPI packages, cumulative binary downloads for the repos that ship them); source-only and self-host apps stay uncounted. A star is a bookmark, not proof anyone runs it.

// DEEP DIVE • READING 9

Right now, one or two people carry each project

Behind the star counts, the **recent** work clusters on very few hands. Across the 200 repos with enough history to judge, over the last six months one author writes a median **44% of the commits**, and **42% had effectively a single committer** in that window.

TOP AUTHOR'S SHARE OF COMMITS, LAST 6 MONTHS • ACROSS 200 REPOS



median project: 44% of the last six months' commits from one author. for 42%, that's one contributor away from stalling its current pace.

EVEN THE FAMOUS RUN ON A FEW RECENT HANDS

TradingAgents	90% one author	93k ★
Claude Code	85% one author	138k ★
Open WebUI	65% one author	145k ★
OpenClaw	49% one author	383k ★


The healthy exception, Immich: 108k stars, and even lately just 8% of commits from its top author, across 900+ lifetime contributors. Proof that big hype and a broad base can coexist. It's just rare.

// THE FIELD, QUANTIFIED 🗒

What it's built in, how it's licensed, who's behind it


The descriptive base under the readings: what the field is built in, how openly it's licensed, and who is actually behind it, by capital and by origin.

📄 LANGUAGE · SHARE OF FIELD



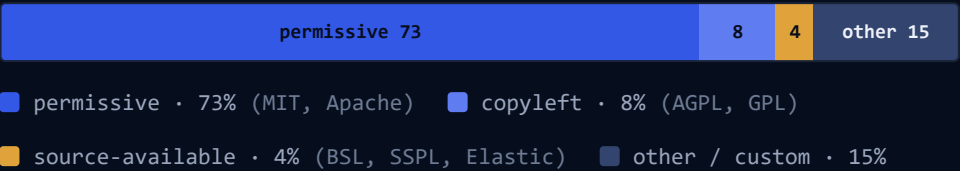
python for the models, typescript for the product, three quarters of the field. rust is the latency layer underneath: 9% and climbing, quietly taking the performance-critical tooling.

💰 WHO'S BEHIND IT · BY CAPITAL



only 10% is foundation-backed: the model with the longest half-life. nearly half (44%) rides on one person or a small unfunded team; VC is already the second force at 29%. the durability question the star count never answers.

📄 LICENSE · IS IT REALLY OPEN?



most of the field is genuinely open. the tell is the 4% that has gone source-available: BSL / SSPL / Elastic / Functional Source, read the code but not run a rival on it. the Terraform / Redis / Elastic relicensing move, now at the frontier: the one shift that can kill adoption overnight.

👤 AND BY ORIGIN · TWO ENGINES

🏢 US big tech

Microsoft 5 · Google 3 · OpenAI 2 · Vercel 2

a land-grab for mindshare

[Playwright MCP](#), [Gemini CLI](#), [Codex](#), [Claude Code](#)

🇨🇳 China-origin labs

HKUDS 7 · ByteDance 2 · OpenBMB 2 · DeepSeek 2

agent tooling at cadence

HKUDS is the most-featured builder overall; [DeerFlow](#), [UI-TARS](#), [Qwen3](#)

// LEADERBOARDS 🦄

What actually moved: four ways

The same field ranked four ways: stars gained (a monthly rate), code shipped, package installs (weekly) and product downloads (total). The top of one board is rarely the top of the next.

⚡ BUZZ • STARS / MONTH	🚀 OUTPUT • COMMITS / WEEK	📦 LIBRARIES • INSTALLS / WEEK	🖨️ PRODUCTS • BINARY DOWNLOADS
1 OpenClaw +62k	1 ELIZA 2,657	1 FastMCP 19M	1 Voicebox 1.2M
2 hermes-agent +57k	2 OpenClaw 2,381	2 BrowserUse 11M	2 OpenCode 1.1M
3 codegraph +43k	3 hermes-agent 761	3 MLflow 10M	3 cc switch 650k
4 headroom +38k	4 hermes-webui 430	4 OpenAI Agents 7.6M	4 onlook 435k
5 taste-skill +23k	5 oh-my-opencode 376	5 Playwright MCP 6.4M	5 RTK 416k
6 OpenCode +20k	6 orca 368	6 OpenClaw 2.1M	6 KittenTTS 230k
7 agency-agents +18k	7 CopilotKit 317	7 chrome-devtools-mcp 2.0M	7 OpenTofu 207k
8 Understand-Anything +17k	8 LiteLLM 314	8 FastAPI-MCP 1.9M	8 Jan 165k
9 RTK +16k	9 codewhale 289	9 Deepeval 1.2M	9 worldmonitor 149k
10 pi-mono +15k	10 Mastra 268	10 FlashInfer 1.1M	10 meetily 127k
attention, not adoption: what the timeline is talking about.	how hard it's shipping: commits merged, week in, week out.	what's actually installed: npm + PyPI pulls per week.	runnable builds pulled from GitHub releases (total).

// WHERE THIS GOES 

Where the vector points next

Fold the nine readings together and one arc runs through them, then three calls we're willing to be wrong about, and the repos our live signals say to watch.

The through-line

Open source's frontier is now **AI, agent-shaped and shipped as products**, often maintained by one or two people, and judged by usage and cost, not stars. The contest moved from "can we build an agent" to "**can we run fleets of them, cheaply, on our own machines.**" And the arrow points one way: AI and agents, coding most of all, are only accelerating, so every signal here matters more each cycle, not less. Step back far enough and this is just open source doing what it does: the abstractions that won this cycle are already turning up in everyone's stack.



Open source is AI's distribution layer

What open source distributes has changed: not just packages developers import, but finished products people actually run, still built and reviewed in the open, community way. The path from launch to a user's machine now runs through a public repo, and the lag is collapsing. Cold start to six-figure adoption fell from years to months.

watch [Open WebUI](#) • [dyad](#)



Value moves to the scaffolding

As models commoditise, the margin shifts to the layers around them: the harness that runs and multiplexes agents, and the cost tooling that keeps them cheap. The next platform contest is fought here, in the open.

watch [headroom](#) • [cc switch](#)



Good models get small enough to run local

The frontier isn't only bigger models. Capable models keep getting smaller, and that's what puts real use cases on local hardware: good-enough intelligence on a laptop or a phone, with no data leaving the machine. That is decisive in privacy-bound verticals (health, legal, gov) the funded-SaaS grid can't easily follow.

watch [Jan](#) • [Unsloth](#)

Written by people who read the commits

Repository Radar is your biweekly, no-fluff scan of the open-source projects worth your attention. Hand-picked by the two of us, every issue, no algorithmic slop.

Alexander Kremer

CO-AUTHOR



Partner at [Keen Venture Partners](#) in Amsterdam, backing AI-native software startups. ~15 years across finance, tech and consulting. Gen-AI and frontier-tech investing at Picus Capital, product and growth at JD.com in Beijing, earlier McKinsey and IBM, and a Schwarzman Scholarship along the way. Occasionally builds the thing instead of funding it.

[LinkedIn](#) ↗

Claudius Seitz

CO-AUTHOR



MIT, CDTM and TUM; now co-founder at [Index Labs](#), building proprietary agentic-AI platforms and custom automation, after investing alongside Alexander at Picus Capital. Reads more diffs than is strictly advisable: roughly the qualification the job requires.

[LinkedIn](#) ↗

FOLLOW THE RADAR 🐦

The biweekly scan, in your inbox

Every two weeks we read the field and send the signal. Live metrics and a get-it block on every repo at [repositoryradar.dev](#).

[Subscribe on Substack](#) ↗