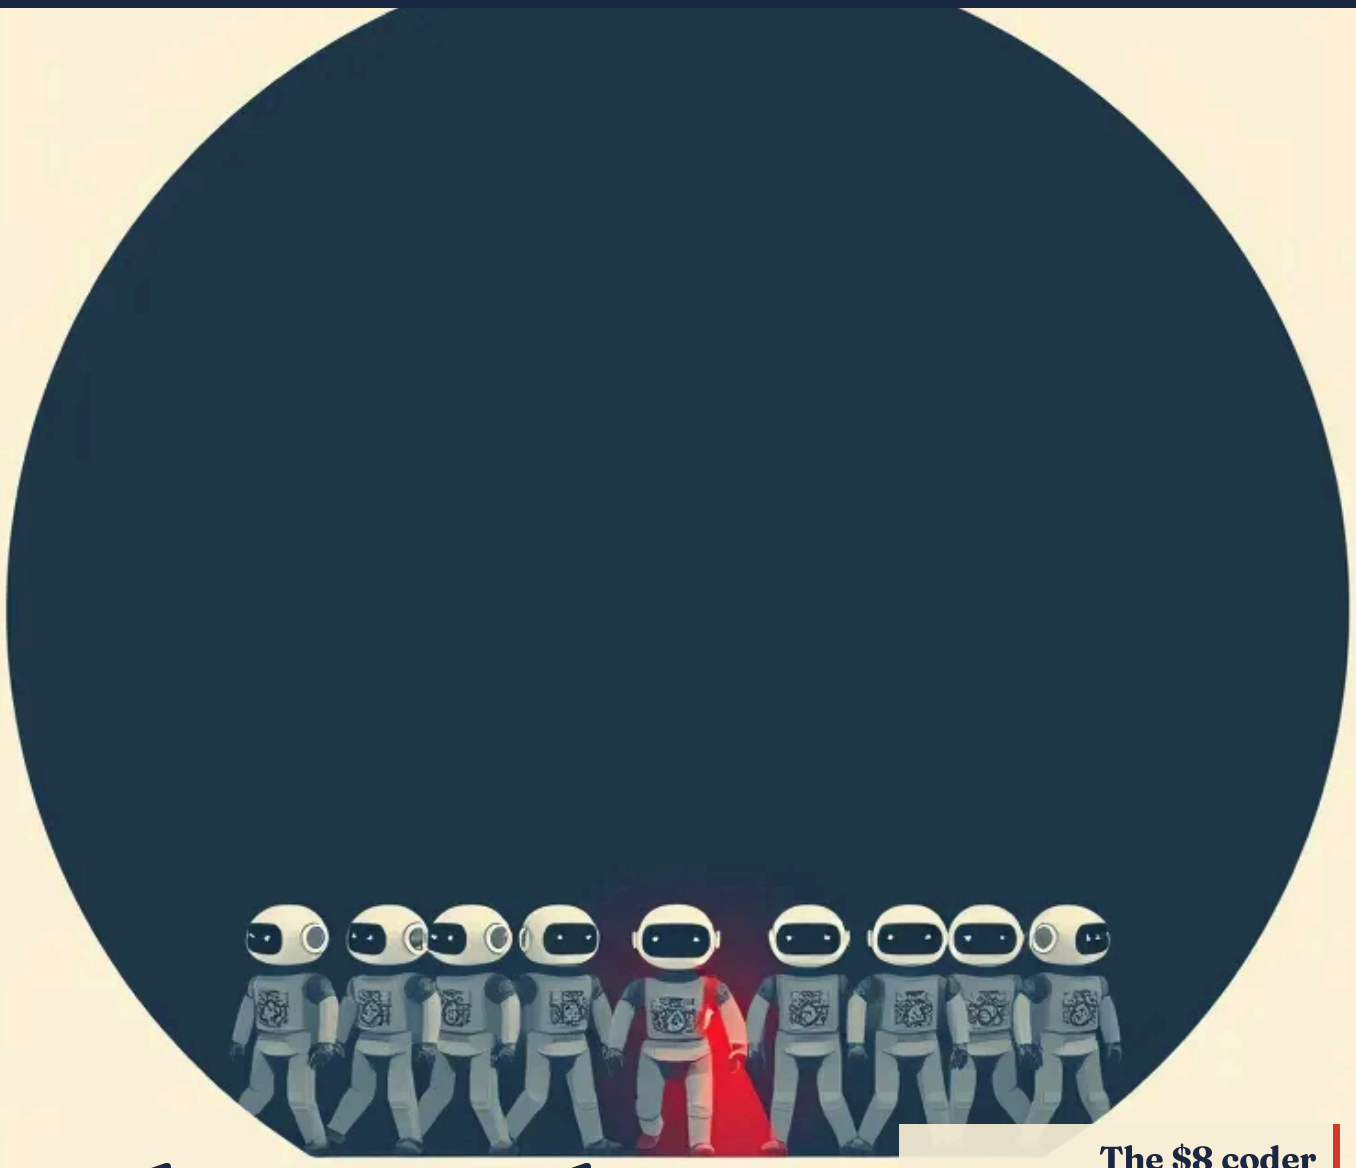
THE AGENT LEDGER

THE MONTH IN AI AGENTS · NO. 0 · OCTOBER 2026 · AGENTGID.COM



Who guards the agents?

A tampered extension, a leaked console and a daemon nobody asked for · p. 4



The $8 coder
What Cursor's new entry plan buys · p. 3

179 incidents
A month of downtime, measured · p. 7

On the rise
Memory, taste, video: what developers star · p. 11

IN THIS ISSUE

Leaders	3
Cover story: Who guards the agents?	4
The Top 20	6
Downtime	7
Prices and limits	8
Releases that mattered	9
Open source	10
Talk of the month	10
On the rise	11
Assets and skills	12
Graphic detail · Retired	13
Agent indicators · Colophon	14

The cover. One robot in the line is not like the others. October's loudest agent stories were about what slips into the queue: a tampered update, a stolen key. Illustration generated with FLUX.1 schnell; type set by AgentGid.

THE MONTH IN ONE PAGE

Oct 7 Anthropic released Claude Haiku 5.5, a small model for high-volume work; the post drew 1,049 points on Hacker News. Two days later Claude Code 2.1.296 made it the default Haiku model.

Oct 9 OpenAI Codex 0.162.1 made GPT-6.1 Sol the default model in its bundled and Amazon Bedrock catalogs (release notes).

Oct 12 Anomaly published OpenCode 2.0: one background service shared by the terminal, web and desktop apps, a rebuilt API and a Code Mode that chains tool calls in one script.

Oct 12 Cursor announced The X Subscription. Its pricing page now starts at Plus, $8 a month; we saw the entry price fall from $20 on Oct 13.

Oct 17 OpenCode said a breach tied to "a large-scale distillation attack" hit 1.4% of OpenCode Console users, those with external model providers linked AS REPORTED (BleepingComputer).

Oct 22 Anthropic donated Agent Skills to the Agentic AI Foundation, the Linux Foundation body that also hosts MCP. It says more than 70 agent platforms have integrated the standard AS REPORTED.

Oct 22 NVIDIA's OpenShell, a sandbox runtime for agents, drew 508 points on Hacker News. Two days earlier it had set a record of 95K weekly downloads.

Oct 26 Anthropic filed an S-1 for a proposed IPO on the Nasdaq Global Select Market under the symbol ANT. Share count and price range are not set.

Oct 29 From about 2:27am to 1:35pm Pacific time, Cursor's extension marketplace served tampered versions of the Claude Code extension that tried to plant an SSH backdoor. Cursor says they "were not developed by Anthropic".

Oct Google rolled out Gemini 4 Argon and, about two weeks later, began pre-release testing of Gemini 4 Carbon for long-horizon engineering.

AGENT SECURITY

Locks before features

October's biggest agent stories were about access, not launches. Buyers should ask three plain questions.

A tampered extension in a marketplace. Provider keys stored in a cloud console. A background process that keeps running after the last window closes. Each of October's loudest stories sat in the gap between an agent and the machine it works on, and none of them was about how well the agent writes code.

The questions for anyone choosing an agent are unglamorous. Where does it keep my keys? Who signs the extensions and updates I install? What keeps running when I stop? A vendor that answers these in its documentation deserves credit in a comparison. A vendor that does not deserves an email.

Attention is already moving that way. NVIDIA's OpenShell, a sandbox runtime, set a record of 95K weekly downloads on October 20 and drew 508 points on Hacker News. Popularity does not prove safety. It does show where developers looked this month.

PRICING

Cheap is a bundle

Cursor's $8 plan is real. Read the plan card before you compare it with anything.

Cursor's new entry plan costs $8 a month. The median entry price among the 14 coding IDE agents with paid plans we track is $19 AGENTGID CALCULATION. On price alone, Cursor went from ordinary to cheapest overnight.

The plan card tells a narrower story. Plus covers a shared usage pool for Cursor and Grok, with "Grok models only". Cloud agents and third-party models at API list price start at Premium, $30 a month. Subscriptions are "only sold directly through X Subscription Corp." Users on the old plans keep their included third-party usage "until you switch", Cursor says.

So $8 buys a different Cursor from the one $20 bought in September. Compare plans by the models and agent features you actually use. A headline price that excludes your model is not a price for you.

BRIEFING

Who guards the agents?

Three incidents in one month show where agent security breaks in practice: in distribution, in stored keys and in background processes.

Coding agents now run with the same rights as the developer who starts them. They read the repository, call the terminal and hold keys to paid model APIs. That makes the path an agent takes onto a laptop as important as the model behind it. October tested that path three times.

The extension

On October 29, between about 2:27am and 1:35pm Pacific time, the Cursor marketplace distributed versions 2.1.312 and 2.1.313 of the Claude Code extension for Windows and macOS. Cursor's security notice says these versions "were not developed by Anthropic" and tried to install a persistent SSH backdoor. The code collected host and user details, added an attacker's SSH key, tried to enable SSH and opened a reverse tunnel.

Cursor removed the versions, took its marketplace offline, rebuilt it and rotated its keys. It says the attacker's server is offline and that the SSH step would likely fail on most machines without admin rights. Phorion, the security firm that reported the backdoor, found it also installed a macOS launch agent named com.apple.sysmond.plist to survive restarts. Its write-up drew 1,013 points and 553 comments on Hacker News.

The console

On October 17 OpenCode told users that a breach "related to a large-scale distillation attack" had affected 1.4% of OpenCode Console users AS REPORTED. Those affected had linked external model providers; the exposed items were stored provider API keys. BleepingComputer reports users seeing unauthorized charges, some above $1,700, and says OpenCode mentioned possible reimbursement.

The timing was awkward. Five days earlier Anomaly had shipped OpenCode 2.0, and the project still climbed four places to No. 3 in our October Top 20.

The daemon

A GitHub issue asking OpenAI to make the Codex CLI background daemon opt-in reached 610 points on Hacker News on October 12. The reporter describes a managed process that keeps running after every Codex window closes, copies about 374 MB of packages into the home folder and checks for updates every hour. When we checked on November 9, the issue showed no reply from an OpenAI maintainer.

Nothing in the issue describes an attack. The complaint is about defaults: software that keeps running and updating itself without asking. Codex 0.163.0, released October 15, added a /daemon command to restart the managed background server.

The response

Developers spent the month adopting fences. OpenShell, NVIDIA's sandbox runtime, set a record of 95K weekly downloads on October 20. E2B, which runs agent code in cloud sandboxes, joined our catalog.

None of this fixes a poisoned marketplace or a leaked key. It does shrink what a compromised agent can reach, which is the part a team controls.

WHAT TO CHECK THIS WEEK

- Used Cursor and installed or updated the Claude Code extension on October 29? Follow Cursor's notice: look for the svc-telemetry SSH key and run its remediation script.
- Linked provider keys to OpenCode Console? Check for OpenCode's email and your provider billing.
- Set a spend limit on every provider key an agent can reach.
- List the background processes your agents start, and turn off the ones you do not need.

Sources: Cursor security notice (Oct 29–30), Phorion (Oct 30–31), BleepingComputer (Oct 19), openai/codex issue 48195, AgentGid data. Hacker News points from the official HN search API, Nov 9.

RANKINGS

Four newcomers, one new No. 3

OpenCode climbed four places to No. 3 and Hermes Agent thirteen to No. 16. DeepSeek Harness and Dify each fell six. Gid Score measures adoption and interest, not output quality.

#	AGENT	MAKER	GID SCORE	MOVE	DOWNLOADS, 7D	GITHUB STARS
1	OpenAI Codex	OpenAI	84	=	27.7M	132K
2	Claude Code	Anthropic	84	▲1	12.1M	154K
3	OpenCode	Anomaly	82	▲4	5M	223K
4	n8n	n8n	80	▲2	160K	209K
5	LangChain	LangChain	79	▼1	41.5M	149K
6	Pi	Earendil	77	▼1	6M	119K
7	OpenClaw	OpenClaw Foundation	76	▼5	3.4M	394K
8	LiteLLM	BerriAI	76	NEW	22.5M	62.4K
9	LangGraph	LangChain	73	▲1	18.8M	44.2K
10	Cursor	Anysphere	72	▼2	—	33.3K
11	GitHub Copilot	GitHub	68	▼2	—	—
12	Langfuse	Langfuse	67	NEW	14M	36.8K
13	AI SDK	Vercel	67	▼2	41.9M	27.6K
14	Firecrawl	Firecrawl	65	NEW	2.9M	197K
15	Claude Agent SDK	Anthropic	64	▲4	25.2M	8.4K
16	Hermes Agent	Nous Research	64	▲13	37.5K	259K
17	Strands Agents	AWS	63	▼5	9M	9K
18	Rea*	morluto	63	NEW	538K	413K
19	Dify	LangGenius	63	▼6	—	160K
20	DeepSeek Harness	DeepSeek	63	▼6	680K	267K

Source: agentgid.com/top/2026-10/. Gid Score and places from the October 31 snapshot, recalculated on November 10 after we stopped counting a parent company's website (such as nvidia.com) toward a product's score. Moves are against the Oct 7 list; New means no place on that list. Downloads (npm and PyPI, 7 days) and GitHub stars as of Oct 31; downloads include CI and automated installs. *Rea carries our "Unusual star growth" label: it gained 415K stars in October, and we have not confirmed the count.

RELIABILITY

179 incidents in four weeks

Vendors posted 179 incidents in our four weekly reports from October 5 to November 1 AGENTGID CALCULATION. The worst week, October 19–25, had 59, 21 of them rated major or critical.

Cursor posted the most incidents in two of the four weeks: 7 in the week of October 5 and 10 in the week of October 26. In the week of October 19, Claude Cowork, Claude Design and Claude in Chrome each logged 7.

The longest incident rated minor or worse came from Testsigma Atto: 44 hours 38 minutes of slowness and intermittent errors in its US region, starting October 23. Gorgias AI Agent had the two longest major incidents, at 27 hours 30 minutes on October 7 and 19 hours 38 minutes on October 30.

A quiet status page is not proof of a quiet month. In the week of October 12, 65 of the 82 agents we watch posted nothing at all. Durations run from the start the vendor's status page gives to the resolution, and the impact rating is the vendor's own. From the week of October 12 an incident a vendor posts for several of its products counts once; the October 5–11 report still counts it once per product.

AGENT	INCIDENT, IN THE VENDOR'S WORDS	RATING	DURATION	STARTED, UTC
Testsigma Atto	Degraded performance in the US region: slowness and intermittent errors	minor	44h 38m	Oct 23, 08:46
Gorgias AI Agent	Aircall V2 not creating tickets	major	27h 30m	Oct 7, 08:57
Gorgias AI Agent	Mobile app authentication issues	major	19h 38m	Oct 30, 20:07
Hebbia	RDS Proxy connection wait high	minor	15h 34m	Oct 9, 22:49
Sourcegraph Cody	Sourcegraph Cloud is degraded in Europe region	minor	10h 4m	Oct 20, 07:38
Cursor	Investigating Service Degradation with VS Code Extensions	minor	9h 34m	Oct 29, 21:38

Sources: AgentGid weekly outage reports for Oct 5–11, Oct 12–18, Oct 19–25, Oct 26–Nov 1, drawn from vendors' official status pages. Incidents without a vendor rating are left out of the table. October 1–4 sits in a report that also covers late September and is not counted.

PRICING

Eleven price changes, one deep cut

Cursor cut its entry price by 60%; most other moves went up. We count a change only when the new price held.

Our four weekly price reports for October list 11 confirmed changes to an agent's entry price or free tier AGENTGID CALCULATION. Switches between monthly and annual billing, readings undone within two weeks and relabeled plans are left out. Warp's Build plan, for example, is $20 a month billed monthly and $18 billed annually. Both are correct; neither is a change.

The biggest move was Cursor's. On October 13 we saw its entry price fall from $20 to $8 a month, and Cursor CLI followed a day later. The new lineup runs from Plus at $8 to Ultra at $200, and Teams starts at $30 a user. See the leader on page 3 for what each tier includes.

Most other moves went up: Juicebox, Kortix, CopilotKit, Elicit, Google Antigravity and OpenHands raised their entry plans. Zapier Agents and Dia cut theirs, and Qodo added a free tier to its credits.

Limits were quiet. Neither Claude nor Codex had a confirmed usage reset in October. On October 31 OpenAI teased "banked resets" for the following days, which our resets tracker lists as a teaser, not a reset.

AGENT	BEFORE	AFTER	SEEN
Cursor	Free + $20/mo	Free + $8/mo	Oct 13
Cursor CLI	Free + $20/mo	Free + $8/mo	Oct 14
Juicebox	Free + $99/mo	Free + $119/mo	Oct 18
Kortix	Free + $40/mo	Free + $50/mo	Oct 20
CopilotKit	Free + $39/mo	Free + $49/mo	Oct 23
Zapier Agents	Free + $33.33/mo	Free + $19.99/mo	Oct 23
Dia	Free + $20/mo	Free + $10/mo	Oct 28
Elicit	Free + $11/mo	Free + $39/mo	Oct 28
Google Antigravity	Free + $19.99/mo	Free + $30.00/mo	Oct 28
OpenHands	Free + $499/mo	Free + $500/mo	Oct 28
Qodo	$0.012/credit	Free + $0.012/credit	Oct 10

Entry price per month, or per credit, as read on each vendor's official pricing page.

Sources: AgentGid weekly price reports for Oct 5, Oct 12, Oct 19 and Oct 26 (1, 3, 3 and 4 changes), AgentGid price watch, cursor.com/pricing (read Nov 9).

CHANGELOG

2,598 releases, eight that change your setup

Tracked agents shipped 2,598 stable releases in four weeks AGENTGID CALCULATION*. These are the ones that change a default, a config key or an upgrade path.*

Oct 7 · Pi 1.1.0	The azure provider was renamed from azure-openai-responses. **Rename the key in your config.**
Oct 9 · Codex 0.162.1	GPT-6.1 Sol became the default model in the bundled and Amazon Bedrock catalogs.
Oct 9 · Claude Code 2.1.296	Claude Haiku 5.5 became the default Haiku model on the Anthropic API.
Oct 11 · OpenCode 1.19.0	Updates are no longer installed automatically; /update installs them or moves you to OpenCode 2.
Oct 15 · Codex 0.163.0	A searchable /config view, follow-up messages that steer a response in progress, and a /daemon command for the background server.
Oct 16 · OpenCode 1.19.1–3	Session sharing was discontinued. Moving to OpenCode 2 is now an explicit choice.
Oct 18 · Claude Code 2.1.299–303	Idle compaction for conversations of 200K+ tokens, new hooks and more OpenTelemetry attributes.
Oct 24 · OpenClaw	Releases 2026.10.6, 2026.10.7 and 2026.9.34 require patched Node: 24.22.1+ or 26.12.1+.

The busiest shipper among the most-used agents was LiteLLM, with 16 releases in the week of October 19 alone.

Sources: AgentGid weekly release reports for Oct 5–11 (479), Oct 12–18 (647), Oct 19–25 (730), Oct 26–Nov 1 (742), from public GitHub releases.

OPEN SOURCE

Documents are the new workload

Four of the ten libraries with the biggest download gains in October read or write slides and PDFs: python-pptx (+94%), PyMuPDF (+41%), PyPDF (+37%) and pdfplumber (+32%). Our reading: agents are being asked to produce office documents, not only code.

Weekly download records

Hermes Agent, 162.7K (up 316%) on Oct 20. OpenShell, 95K on Oct 20. MCP for Unity, 76.5K (109% above its average) on Oct 14.

New on our shelves

LiteLLM, Langfuse, E2B, Context7, Composio, Mem0, Firecrawl and CopilotKit were among 20 agents added to the catalog. A new entry does not mean a new product.

TALK OF THE MONTH

What developers argued about

STORY, AS POSTED ON HACKER NEWS	POINTS	COMMENTS
Claude Haiku 5.5 · Oct 7	1,049	489
Backdoor in Claude Code for VS Code extension on the Cursor Marketplace · Oct 30	1,013	553
OpenCode Hacked: User credentials leaked · Oct 17	977	185
OpenCode 2.0 · Oct 14	890	397
Codex CLI installs a persistent, self-updating background daemon by default · Oct 12	610	294
Claude Code regularly sends each background agent's transcript to its model · Oct 29	591	236
Claude Duet · Oct 21	568	539
OpenShell – A safe, private runtime for fleets of autonomous AI agents · Oct 22	508	217
How the Claude Code team does code review · Oct 16	488	481
AI coding agents generate more code, but not more software · Oct 12	365	314

Points and comments from the official Hacker News search API, read Nov 9. Dates are HN posting dates, UTC.

Sources: AgentGid Top of October 2026 (library gains, new entries), AgentGid record reports.

ON THE RISE

Four ideas developers are betting on

The fastest-growing projects of October cluster around four jobs: giving agents a memory, giving them taste, letting them make video, and running several of them at once. Each entry says what the project does, which problem it solves and why it is climbing.

1 · MEMORY

Agents that remember

The problem. A coding agent starts every session from zero. Memory tools record what happened, compress it into summaries and feed the useful parts back next time.

Who is climbing. Claude Mem gained 13.6K GitHub stars in October; its npm downloads reached 35.2K a week, up 61% over 30 days. Hindsight from Vectorize (MIT) gained 13.5K and works as a library, an API or an MCP server. ECC (+25.5K) bundles memory with skills, hooks and rules so the agent plans, tests and reviews its own work.

IDEA TO BUILD

Memory a team can share and audit: what the agent "knows" about a repository, who taught it, and a way to make it forget.

2 · TASTE

Skills against the generic look

The problem. Pages built by agents tend to look alike. Taste skills give the agent an art direction and check the result before it reports back.

Who is climbing. Hallmark (MIT, for Claude Code, Cursor and Codex) gained 14.9K stars. Anatomy (+23.4K) draws interactive isometric SVG illustrations of machines and systems. Taste Skill picks one of 20 art directions and checks the page in a headless browser at three screen widths.

IDEA TO BUILD

A house-style skill generated from a company's own site and design system, kept in its repository and versioned like code.

3 · VIDEO

Video written as code

The problem. Turning a brief, a pull request or a dataset into a video without a timeline editor. When the scene is HTML or code, an agent can write it and the render repeats exactly.

Who is climbing. HeyGen's HyperFrames (Apache-2.0) gained 22.1K stars, and its npm downloads rose 90% to 3.6M in October. Its README lists launch videos, PR walkthroughs and chart races. OpenMontage (+12.4K, AGPL-3.0) turns a plain-language brief into a finished video and shows the cost of each asset before rendering.

IDEA TO BUILD

Release notes that become a 30-second video on every tag, so a changelog channel runs itself.

4 · PARALLEL AGENTS

Many agents, one desk

The problem. Developers now run several coding agents at once. They want to compare the results, switch models and keep the bill in view.

Who is climbing. Orca (MIT) had the month's biggest gain, 29.9K stars: Claude Code, Codex, OpenCode and Pi side by side, each in its own git worktree, with a phone companion. Magpie (+23.2K) chooses the model behind each agent, such as Codex on DeepSeek.

IDEA TO BUILD

A referee for parallel runs: one task goes to three agents, the tests decide which worktree is merged, and each attempt's cost is logged.

Also climbing: Moli (+19.2K), a headless browser written in Rust that agents use to read the web, and Graphify (+18.1K), which maps a codebase into a graph that agents query instead of searching files.

Star gains for October from agentgid.com/top/2026-10/, read Nov 9; descriptions from AgentGid cards and the projects' GitHub pages. Download counts include automated installs. "Idea to build" is the editors' suggestion, not a product.

ASSETS

Voice and video lead the gains

Three of the five open assets that grew fastest in October render video; two make speech. All five run locally.

ASSET · WHAT IT DOES	OCT DOWNLOADS	CHANGE
Chatterbox · text-to-speech with voice cloning, 23 languages MIT · every file carries a watermark	243K	+97%
HyperFrames · renders HTML, CSS and animation to MP4 Apache-2.0	3.6M	+90%
Revideo · TypeScript scenes rendered to MP4 MIT	121K	+77%
Manim · math explainer animations in Python MIT	211K	+71%
Piper · fast local text-to-speech on CPU GPL-3.0 · each voice has its own licence	2.1M	+71%

Change is October against September, for assets with at least 20K downloads in September. Licence notes from each AgentGid card; check the licence page before you ship.

SKILLS

The most-starred packs teach habits

The two biggest skill packs teach process rather than knowledge: plan first, test first, review before finishing. The next wave teaches taste.

SKILL PACK · WHAT IT TEACHES THE AGENT	STARS
mattpocock/skills · interview the plan, write a spec and small tickets, build test-first 39 skills · MIT	319K
obra/superpowers · clarify, get a design approved, plan, build test-first, review 15 skills · MIT	315K
anthropics/skills · Office files and PDFs, design, web app testing, writing new skills 19 skills · mixed licences	190K
Leonxlnx/taste-skill · art direction for pages, checked in a browser at three widths 1 skill · MIT	112K
nexu-io/open-design · pages, decks and reports in a brand's style, taken from its live site Apache-2.0	107K

Star totals from agentgid.com/skills/, read Nov 9. From the November issue this table ranks packs by the month's star growth.

HOW A PROJECT GETS INTO "ON THE RISE"

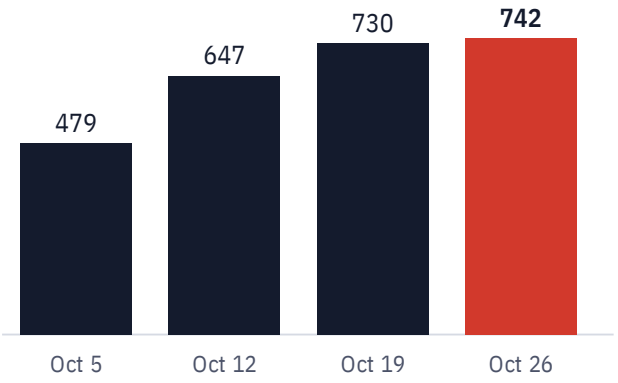
- It is in the AgentGid catalog and has an official repository or homepage.
- It gained the most GitHub stars in the month, or its downloads grew fastest from a base of at least 20K the month before.
- Outliers are held back until checked. In October one project showed a gain of 415K stars in a month; it is left out.
- Each entry says what the project does and which problem it solves, in plain words, with a link to its card.

GRAPHIC DETAIL

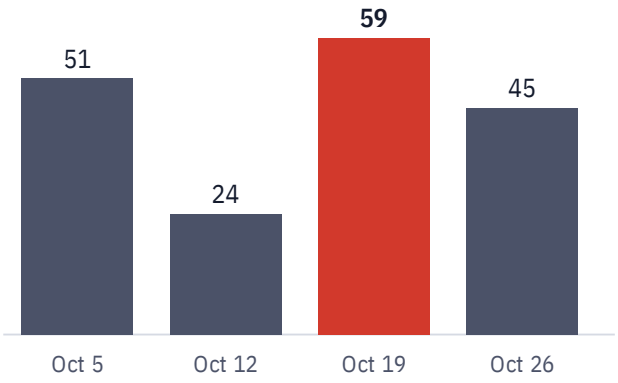
Shipping sped up every week; outages did not follow

Stable releases rose 55% from the first to the last week of October. Incidents swung between 24 and 59 with no trend AGENTGID CALCULATION.

Stable releases per week



Incidents posted per week



Weeks start on the date shown. Each panel has its own scale. Source: AgentGid weekly reports.

RETIRED THIS MONTH

- **OpenCode session sharing** was discontinued in 1.19.x (Oct 16).
- **Automatic updates in OpenCode 1.x** stopped with 1.19.0 (Oct 11).
- **Cursor's old Ultra plan** is no longer available for new signups, according to a reply in Cursor's X Subscription thread.

AGENT INDICATORS

October in one ledger

INDICATOR	OCTOBER
Agents tracked (Oct 31)	480
Status pages watched	82
Incidents posted, Oct 5–Nov 1	179
Stable releases, Oct 5–Nov 1	2,598
Confirmed price changes, Oct 5–Nov 1	11
Median entry price, coding IDE agents	$19
Downloads across tracked agents, week 43	202.3M
Highest Gid Score (OpenAI Codex)	84
New entries in the Top 20	4

Incident, release and price-change totals add up our four weekly reports AGENTGID CALCULATION.

Month-on-month columns start with the November issue. Downloads: week 43 index.

COLOPHON

How this issue was made

The Agent Ledger is AgentGid's monthly review of AI agents. Most figures come from AgentGid's own data: the catalog, status checks every 30 minutes, monthly pricing snapshots and GitHub releases. The rest come from official pages: vendor blogs, pricing pages, security notices and the Hacker News API. Every figure links to its source.

Vendor claims are marked AS REPORTED. Our own arithmetic is marked AGENTGID CALCULATION. Gid Score measures adoption and interest, not quality. We do not test products for this magazine.

Texts are drafted by an AI editor from a frozen snapshot of the month, checked line by line against the sources by a second, independent checker, and approved by a human editor before publication. The cover illustration is generated with FLUX.1 schnell; the masthead and type are set by AgentGid.

Found a mistake? Corrections appear in the next issue and in this one, marked with the date.