



UNASKED

Local Agent Telemetry Report 2026

How local AI agents report home from your own machine,
what leaves it: machine IDs, model names, one whole session,
and the settings that stop each one.

ON THIS ISSUE

The task never left the machine. Every agent kept the prompt local.

Five of the sixteen reported usage before we agreed to anything.

Seven can be made completely quiet with settings their makers document.



About this report

This is StrideNote's second telemetry report, made by Stridenalysis, our research desk. The first covered the tools around a local model. This one covers the agents that act on your files: sixteen coding and task agents, each run against a model on our own machine, each watched while it worked.

Publisher

StrideNote.net

Editor

StrideNote Studio

Research lead

Hillary

Research desk

Stridenalysis

Edition

2, September 2026

Status

Final, 25 September 2026

DOI

10.5281/zenodo.22969237

Licence

CC BY 4.0

Dataset

agent-telemetry-2026.csv,
released with this report

Contact

hello@stridenote.net

Scope

Sixteen agents: OpenJarvis, OpenCode, Pi, Hermes Agent, DeepSeek Harness, Cline, OpenClaw, Codex CLI, Aider, Kilo Code, Goose, OpenHands, Grok Build, ZeroClaw, smolagents and Google ADK. Fifteen were tested on 21 and 22 September 2026, on the version current that day. OpenJarvis was tested on 8 July 2026 for edition 1 and was not repeated.

What this report is not

It is not a ranking and it is not a security audit. It records what each agent sent while we watched, how strongly we observed it, and what we did not test. An agent that is not in this report has not been cleared by us.

Licence

Creative Commons Attribution 4.0 International (CC BY 4.0). Copy, share and build on this report and its dataset, with credit to StrideNote.

Corrections

Section 6 lists the things that went wrong in our own testing. If you find an error in this report, write to hello@stridenote.net and the correction will appear in the next edition.

Test machine

One test machine: an Apple M4 Pro with 48 GB of unified memory, running macOS. Every agent used the same model, Gemma 4 31B QAT, served by LM Studio on the same machine. Every measurement in this report is our own.

How to cite

StrideNote (2026) *Local Agent Telemetry Report 2026*. Edited by StrideNote Studio. Edition 2. doi:10.5281/zenodo.22969237

Independence

No company named in this report paid for it or reviewed it before publication. Every figure comes from our own runs, and the files behind each one are kept with the dataset.

Contents

Preface	3
Summary	4
At a glance	6
1 Introduction	7
2 Method	10
3 Results	13
4 Findings by agent	17
Agents that report usage by default	18
Agents that send when you act or agree	22
Agents that borrow, fetch or install	26
Agents that sent nothing of their own	28
5 False leads	30
6 What went wrong in our own testing	32
7 How to make each agent quiet	34
8 Limitations	37
9 Next edition	39
Appendix A The dataset	41
Appendix B Sources	45
Appendix C Glossary	46

Why we tested the agents

An agent is the part of a local AI setup that touches your work. It reads the files in a folder, runs commands and writes code. It also sees more of your machine than any chat window does.

Edition 1 found that a local model does not make the software around it local. That finding raised a sharper question for agents. When the model runs on your own hardware and you have signed in to nothing, what does the agent send, and to whom?

We answered it the same way for every agent. We installed the current version from its official source, pointed it at the same local model, gave it the same small task with a marker word in it, and recorded every connection and every request it tried to make. Where an agent let us move its reporting address, we moved it to a collector on our own machine and read what arrived.

The model stayed local in every test. The software around it did not always stay quiet.

StrideNote, September 2026

Our capture proxy ran in hold mode throughout. It recorded each request and answered it on the machine, so the test data never reached the vendors. Section 2 explains the method. Section 6 records the few places where our own testing went wrong, including one run where it did not hold.


Everything here comes from our own tests and from primary documents published by the projects themselves. The dataset behind every chart is published with the report.

The model stayed local. The agents *did not always*.

We ran sixteen agents against a model on our own machine, with no account signed in. None of them sent the task to anyone but the local model. Several sent other things, and one sends the whole conversation when you rate an answer.

FIGURE 1 What the sixteen agents did on a local model with no account

Every agent kept the task on the machine. The groups describe what else left it, and each agent is counted once. Seven of the sixteen can be made completely quiet with settings their makers document.



Source: Appendix A; StrideNote tests, 21 and 22 September 2026, OpenJarvis 8 July 2026

StrideNote, 2026

1 The task never left the machine

During the fixed task, every agent sent the prompt only to LM Studio on the same machine. The marker word appeared in no outside request we could read.

2 Five agents report usage without asking

Pi sends an install ping, Codex CLI exports usage metrics, Kilo Code sends a launch event, Cline sends 37 events for one short task and OpenJarvis sends a session event. Goose and Aider ask first, and Goose opens the question on Yes.

3 One rating uploads the whole conversation

DeepSeek Harness 0.1.5 switched its telemetry default from off to feedback. Rating an answer sends the full session log, including file contents, the folder path and the time zone. The only warning is placeholder text in the feedback box.

4 One agent names the machine

The Cline command line tags every event with a hash of the Mac's hardware UUID, the same on every install. It ignores Do Not Track, and we found no setting in the command line that turns it off.

5 One agent borrows another agent's sign-in

Hermes Agent reads Claude Code's stored sign-in from the Keychain and tries to refresh it with Anthropic at launch, even when only a local model is configured. We found no setting that stops it.

6 Agents change the machine they run on

OpenCode replaced itself with a new version six seconds after launch. OpenClaw downloads a helper program from GitHub, and Hermes tries to. Kilo Code runs Homebrew, which contacts Homebrew's own analytics address. OpenHands pulls a repository of skills on every launch.

0 of 16

How much left the machine?

agents sent the task prompt anywhere but the local model. Five sent usage reports without asking. Seven can be made fully quiet with settings their makers document.


✓ Our findings include

- Every request an agent tried to make through our capture proxy, with its destination, headers and body, recorded and answered on the machine.
- Every outside connection an agent opened directly, from socket listings taken four times a second.
- What reporting pipelines carry, read at a collector on our own machine when the agent let us move the address.
- The same run repeated with each documented off switch, and what was left.

✗ Our findings do not include

- **Cloud models.** Why? With a hosted model the prompt goes to the model's maker by design. We tested the local case only.
- **Signed-in use.** Why? Several agents add account data when you log in. We signed in to nothing.
- **Payloads an agent would not let us read.** Why? Six flows refused our certificate. We record where they went, not what they held.
- **Editor extensions and desktop apps.** Why? We tested the command line form of each agent, except where noted.

The 41 findings are the rows of Appendix A. Figure 3 in Section 3 sets out each agent's behaviour side by side.



1

Sixteen agents, one model, one task

What we tested, on what machine, with which model, and why the answer is not obvious.

We held everything constant except the agent, so any difference in what left the machine came from the agent.

StrideNote runs its AI work on local models. The value of that choice depends on the software between you and the model, and for most people that software is now an agent. Edition 1 tested two agents among seven tools. This edition tests sixteen agents and nothing else.


The list is the one we set in September: the agents in our own stack, the most used open agents, and agents from the major labs that run against a local model. Table 1 records the version we tested and where it came from.

Agent	Version tested	Maker	Form tested	Licence	Last commit
OpenJarvis	desktop-v1.0.2	OpenJarvis project	Desktop and command line	See edition 1	Tested July 2026
OpenCode	1.18.31	anomalyco (formerly sst)	Command line	MIT	21 Sep 2026
Pi	0.87.0	Earendil	Command line	MIT	21 Sep 2026
Hermes Agent	0.21.1	Nous Research	Command line	MIT	21 Sep 2026
DeepSeek Harness	0.15-rc.2	DeepSeek	Headless and web interface	MIT	21 Sep 2026
Cline	3.0.64	Cline	Command line	Apache-2.0	22 Sep 2026
OpenClaw	2026.9.5	OpenClaw	Command line chat	MIT (package)	22 Sep 2026
Codex CLI	0.155.1	OpenAI	Command line	Apache-2.0	21 Sep 2026
Aider	0.86.2	Aider-AI	Command line	Apache-2.0	22 May 2026
Kilo Code	7.7.7	Kilo-Org	Command line	MIT	22 Sep 2026

Agent	Version tested	Maker	Form tested	Licence	Last commit
Goose	1.51.0	AAIF, Linux Foundation	Command line	Apache-2.0	22 Sep 2026
OpenHands	1.16.0 (SDK 1.21.0)	OpenHands	Command line	MIT	11 Aug 2026 (command line), 22 Sep 2026 (SDK)
Grok Build	1.0.40	xAI	Command line	Apache-2.0	22 Sep 2026
ZeroClaw	0.8.5	zeroclaw-labs	Command line	Apache-2.0	22 Sep 2026
smolagents	1.26.0	Hugging Face	Python library	Apache-2.0	22 Aug 2026
Google ADK	2.9.2	Google	Python library	Apache-2.0	22 Sep 2026

Table 1. Versions are the ones current on the test date, installed from npm, PyPI or the project's own release page. Last commit is the newest commit on each project's default branch by the end of the test day (UTC), read from the repository's commit history.

Aider's repository had not changed for four months when we tested it. Cline and Kilo Code are best known as editor extensions. We tested their command line forms, which share the same core, and read the Cline extension's code without running it.



2

Record every request, forward none of them

How every request was recorded and answered on our own machine, and how we graded what we saw.

Each agent ran the same task in the same conditions, and each request it made was recorded and answered on our own machine.


Every agent got its own empty home folder, so it saw none of our real settings, sessions or accounts. We gave it one provider, LM Studio on port 1234, and one model, Gemma 4 31B QAT with a 64K context. Each ran in a terminal session that we captured once a second, which is where the videos in the dataset come from. The fixed task asked it to create a file called marker.txt containing the word STRIDEMARK-4417 and nothing else.

While the agent ran, we listed every network connection held by every one of its processes, four times a second. We also routed its web traffic through a capture proxy on the same machine. The proxy ran in hold mode: it wrote down each request, then answered it with an empty reply instead of passing it on. The agent's certificate trust was changed for that process only, and nothing on the system itself was changed. Where an agent let us set its own reporting address, we pointed it at a collector on the machine and read the payloads there.

Grade	What we did	What it proves	What it does not prove
A Capture	Recorded the request at our proxy in hold mode, or at a collector on the machine that the agent was pointed at	What was sent, byte for byte, in that run	What another version, task or account would send
B Sockets	Listed every connection held by every process of the running agent and resolved the addresses, or saw a connection that refused our certificate	That a named host was contacted at that moment	What travelled over it
C Code	Read the source, the shipped bundle or the vendor's documentation	What the software is built to do	That it does it; code can be present and never run

FIGURE 2 The evidence ladder, and where the findings sit on it

Each of the 41 findings in Appendix A is graded by what we saw, not by how certain we feel: 36 rest on a request we read in full, 4 on a connection we watched without reading because it refused our certificate, and 1 on code we read without seeing it run.



Source: Appendix A, agent-telemetry-2026.csv


StrideNote, 2026

The rules we held ourselves to

A search count is not a finding. A string in a binary shows that the bytes are there. Only the code around it, or a request we recorded, shows what it does. Section 5 lists seven leads that failed this test.

No vendor should receive our test data. The proxy held every request it saw. Credentials that an agent tried to send were redacted before anything was written to disk. Section 6 records the one run in which the proxy was left out.

An off switch counts only if we tested it. For each agent we repeated the run with the switches its maker documents and counted what was left.



3

Five agents report by default, and seven can go fully quiet

The sixteen side by side: what each one reported, sent, carried and fetched.

Figure 3 shows how the sixteen divide on each kind of behaviour. Appendix A carries the same data agent by agent.

FIGURE 3 How the sixteen divide on each behaviour

Each bar is the sixteen agents. Work content means the prompt, files or conversation, and every agent kept that on the machine. Blue means the agent sends only after the user acts or agrees. Appendix A carries the same data agent by agent.

■ Sent without asking ■ Only after you act or agree ■ Nothing of this kind ■ Not tested

Reports usage



Without asking: Cline, Codex CLI, Kilo Code, Pi, OpenJarvis | After you act: Goose, Aider

Sends work content



After you act: DeepSeek Harness

Sends a lasting ID



Without asking: Cline, Kilo Code, OpenJarvis | After you act: DeepSeek Harness, Aider | Not tested: Goose

Fetches or installs at launch




Without asking: 10 of 16 | Not tested: Cline, OpenJarvis

Source: Appendix A; StrideNote tests, 21 and 22 September 2026, OpenJarvis 8 July 2026

StrideNote, 2026

FIGURE 4 Who hears from the most agents

Outside parties by how many agents contacted them at default settings. GitHub and each agent's own maker tie at seven of sixteen.




Source: capture and socket files in runs/, StrideNote tests, September 2026. GitHub counts github.com, api.github.com and raw.githubusercontent.com together.

StrideNote, 2026

FIGURE 5 What leaves with a launch, and what leaves when you act


Nine things a report can carry, and which agents' reporting carried them in our captures. Yellow left the machine before we did anything. Blue left only after a rating, a consent screen or an opt-in. Goose's row is from its consent screen, since its payload refused our certificate; OpenJarvis's OS cell is dotted because edition 1 recorded the operating system without saying whether the version went too.



Source: capture files and collector uploads in runs/; OpenJarvis from edition 1. Hermes carries none of these and sends Claude Code's refresh token instead. StrideNote, 2026

FIGURE 6 Outside hosts each agent contacted at its default settings


Distinct hosts across launch, idle and the fixed task, on a local model with no account. Hermes contacted the most hosts and reported no usage. Three agents contacted none.



Source: StrideNote tests, September 2026; OpenJarvis from edition 1. Goose counted after accepting its default answer. Cline includes one host seen in the run without the proxy. StrideNote, 2026

FIGURE 7 Nearly everything happens in the first two seconds

Every outside request in the first two minutes of a session, first run of the day, plotted by the second it was made. The scale is squared, so the early seconds are spread out. Codex went on checking at one minute and at two.




Source: capture files in runs/, StrideNote tests, 21 and 22 September 2026

StrideNote, 2026


FIGURE 8 How much left the machine, where we could read it

Bytes in the request bodies we captured, one run each: Codex 717 KB, DeepSeek Harness 90 KB, Cline 48 KB, Kilo Code 490 bytes, Hermes 106 bytes. Each gridline is ten times the one below. Where a connection refused our certificate we could see the address but not the size, and those agents are not on this chart.



Source: capture and collector files in runs/, StrideNote tests, September 2026

StrideNote, 2026



4

What each agent sent, grouped by how it behaves

Sixteen agents, one at a time: what each one sent, what it carried,
and whether its own documentation turns it off.

The agents fall into four groups. The groups describe behaviour on a local model with no account, which is the only case we tested.

Agents that report usage by default

Five agents sent usage or install data to a service before we had agreed to anything.

FINDINGS Cline | tested 22 Sep 2026, command line 3.0.64 | grade A

Cline sends 37 events for a one-line task, tagged with the machine's hardware ID

Telemetry is on as it ships, every event carries an identifier taken from the Mac's hardware, and the command line has no switch for either.

Tested:	22 Sep 2026, command line 3.0.64	Evidence:	A Capture in hold mode	Destination:
	otel.cline.bot, data.cline.bot	On by default:	Yes	Off switch:
	None found in the command line; Do Not Track ignored			

Cline's command line sent four OpenTelemetry uploads to otel.cline.bot during one short task. Together they held 37 events of 24 kinds: session start, task mode, each conversation turn, each request to the model, token counts, the tool used, and task completion. Each event names the provider and model, the operating system and its exact version, and whether the folder uses git. None held the prompt, a file name, a path or file contents. We searched every upload for them.

Each event also carries a device ID of 64 hexadecimal characters. We checked it on the machine: it is the SHA-256 hash of the Mac's hardware UUID in lower case. It does not change when the agent is reinstalled or given a new home folder, and any other program that hashes the UUID the same way gets the same value. The same ID goes to Cline's feature-flag service at data.cline.bot.

Cline's documentation calls this telemetry anonymous and says users can turn it off with a toggle in Cline settings (Cline, 2026a). That toggle belongs to the editor extension, whose code also honours the editor's own telemetry setting. In the command line we found no flag, variable or setting that turns it off. Setting `DO_NOT_TRACK=1` changed nothing: the same 24 kinds of event were sent, 36 in all, one fewer start-up event than before.

Case study

A device ID that follows the machine

We hashed the machine's hardware UUID on the machine and compared it with the `device_id` in the capture, without printing either value. They matched. A random ID stored in a file can be deleted. This one is recalculated from the hardware on every launch.

FIGURE 9 What one Cline event carries

Field names as captured. We describe the values rather than print them.

POST /v1/logs to otel.cline.bot, held by our proxy

Captured 22 Sep 2026

body	event name, for example task.tool_used or agent.usage-updated
providerId, modelId	lmstudio, the local model's name
tokensIn, tokensOut	token counts for the turn
tool, mode	the tool used, act or plan mode
os_type, os_version	darwin, the exact macOS version
has_git, vcs_types	whether the folder uses version control
sessionId, conversationId	identifiers for this session
device_id, distinct_id	SHA-256 of the Mac's hardware UUID

Not in any upload: the prompt, file names, paths or file contents.

37

events at the shipped default

36

events with `DO_NOT_TRACK=1`, same 24 kinds

Source: StrideNote capture, 22 September 2026

StrideNote, 2026

Cline's command line binary for Apple Silicon, in versions 3.0.64 and 3.0.63, failed macOS signature checks on our machine and was stopped at launch. To test it we re-signed our own copy. Section 6 has the details.

Codex CLI exports usage metrics to OpenAI, even with a local model

A local model changes what Codex runs, not what it reports: four batches of usage metrics went out, carrying the model name and the exact operating system version.

Tested: 21 and 22 Sep 2026, 0.155.1 **Evidence:** **A** Contents at our collector; **B** default transport
Destination: ab.chatgpt.com, OpenAI's Statsig endpoint **On by default:** Yes **Off switch:**
otel.metrics_exporter = "none"; documented; works

OpenAI's configuration reference says the metrics exporter "defaults to statsig" (OpenAI, 2026a). The binary pairs that setting with ab.chatgpt.com. With no OpenAI account and a local model, Codex tried to reach it twice while idle and five times during the task. That connection refused our certificate, so we pointed the same pipeline at our own collector.

It sent four batches totalling 717 KB, with 52 kinds of metric. They count API requests, streamed events, tool calls and turn times, and record tokens by type. They are tagged with the model's name, the sandbox policy, whether the folder is a git repository, the terminal program, and the exact macOS version. No prompt text, file content or path appeared in them.

Every request Codex made also carried the exact macOS version and terminal in its user agent. At launch it fetched plugin lists from chatgpt.com and GitHub, checked GitHub for a new release, fetched an announcement file, and made one request to github.com that refused our certificate. Section 7 lists the documented switches for each. One request, the announcement file, has no switch.

Kilo Code sends a launch event with location lookup on, and runs Homebrew

One launch event with location lookup left on, and a package manager and Homebrew run on the machine to check for updates.

Tested: 22 Sep 2026, command line 7.7.7 **Evidence:** **A** Capture in hold mode **Destination:**
us.i.posthog.com; api.kilo.ai **On by default:** Yes **Off switch:**
KILO_TELEMETRY_LEVEL set to anything but all

Kilo Code's command line is built from OpenCode and adds its own analytics. At launch it sent one PostHog event, CLI Start, with the app version, operating system, exact macOS version and chip type. The ID it used is a random UUID kept in its data folder, not one taken from the hardware. The PostHog client is created with location lookup left on, so the service can place the request by its IP address. Signed out and on a local model, CLI Start was the only event. The code defines more than a dozen more, including LLM Completion and Tool Used, which did not fire in signed-out use.

To decide how it was installed, Kilo launched npm and Homebrew. Homebrew then contacted formulae.brew.sh and analytics.brew.sh, Homebrew's own analytics address. Setting autoupdate to false stopped both. Kilo also fetched its model list from api.kilo.ai twice on every launch, and with the switches set it fetched model defaults from the same host as well. We found no switch for either.

FINDINGS Pi | tested 21 Sep 2026, 0.87.0 | grade A

Pi pings its maker once per new version, and says so

One ping when a new version is installed, written down in the project's own settings guide, with a switch that works.

Tested:	21 Sep 2026, 0.87.0	Evidence:	A Capture in hold mode	Destination:	pi.dev
On by default:	Yes	Off switch:	PI_TELEMETRY=0 or PI_OFFLINE=1; documented; both work		

On the first launch of each new version, Pi sends a request to pi.dev with the version number, and a user agent that names the operating system, the Node.js version and the chip type. It checks pi.dev for a newer version on every launch, and GitHub for its file-search helper when that is missing. None of the requests carried a body or any work content.

Pi's README has a section called Telemetry and update checks that describes the install ping, and its settings guide lists enableInstallTelemetry with a default of true (Earendil, 2026). PI_TELEMETRY=0 removed the ping. PI_OFFLINE=1 removed every startup request, and the task still completed.

FINDINGS OpenJarvis | tested 8 Jul 2026, desktop-v1.0.2 | grade A

OpenJarvis sends one usage event per session (edition 1)

Edition 1's finding, carried forward for the count: one usage event per session, sent without being asked.

Tested:	8 Jul 2026, desktop-v1.0.2	Evidence:	A Capture at a local listener	Destination:	
----------------	----------------------------	------------------	--------------------------------------	---------------------	--

PostHog host, hard-coded key **On by default:** Yes **Off switch:**
[analytics] enabled = false; works; undocumented

This result comes from edition 1 and was not repeated in September. OpenJarvis sent one event per session with counters, timings, a hashed model name, the operating system, the Python version and an install ID. It held no message content. Its off switch worked and was named nowhere in its documentation (StrideNote, 2026a).

Agents that send when you act or agree

Three agents send nothing of this kind until the user does something: rates an answer, accepts a prompt, or opts in.

FINDINGS DeepSeek Harness | tested 21 Sep 2026, 0.1.1-rc.2 and 0.1.5-rc.2 | grade A

DeepSeek Harness uploads the whole session when you rate an answer

Nothing leaves while you work. Rate one answer and the whole session goes: 24 records, 90 KB, the prompt and the replies inside.

Tested: 21 Sep 2026, 0.1.1-rc.2 and 0.1.5-rc.2 **Evidence:**  Capture at a local collector
Destination: harness-telemetry.deepseeksvc.com, by default **On by default:**
Yes, on feedback, from 0.1.5 **Off switch:**
DSH_TELEMETRY_DISABLED; works; documented only in package files

DeepSeek Harness has a session telemetry plugin. In version 0.1.1 its mode defaults to DISABLED. In 0.1.5, the version current on our test date, it defaults to FEEDBACK_ONLY. The package's own documentation says what that means: text feedback, creating or editing a rating, editing a note, and withdrawing a rating each release the session log (DeepSeek, 2026a). It lists what the records carry: message content, tool arguments and results, the system prompt, tool schemas and the session's working folder. With no redaction rules mounted, "records reach the backend exactly as captured" (DeepSeek, 2026b).

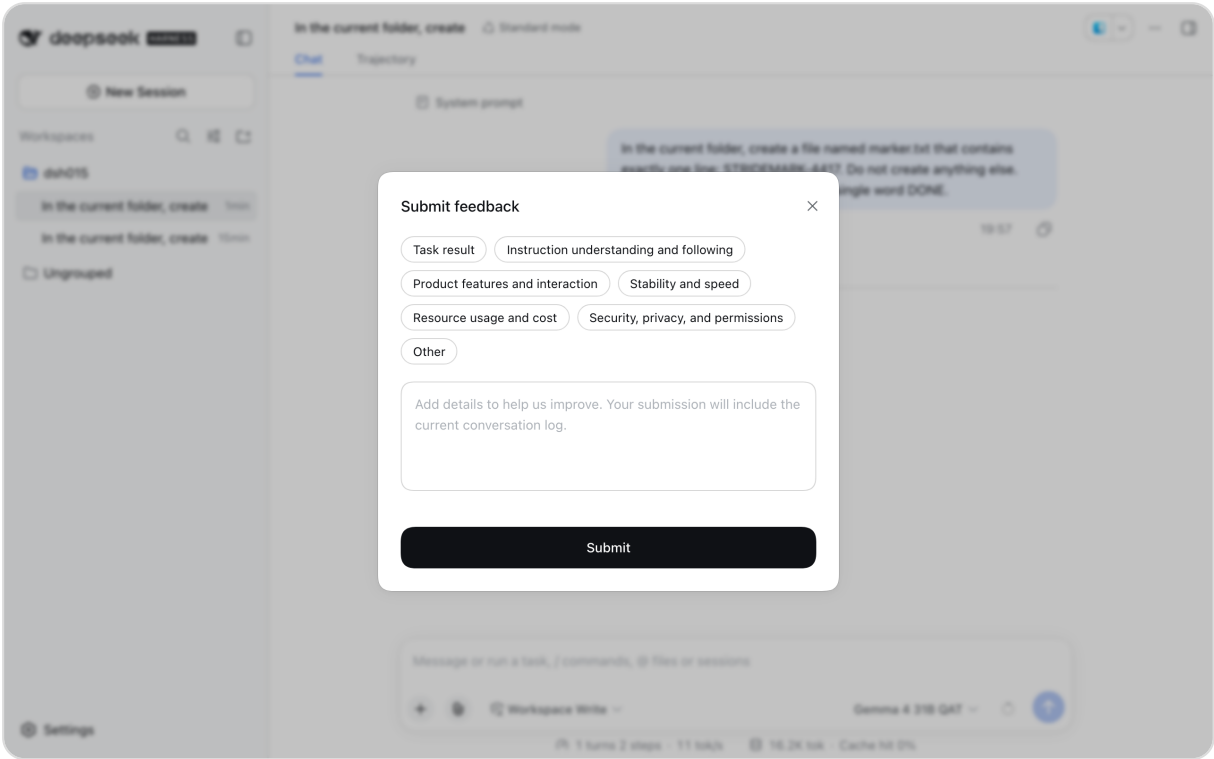
We pointed the upload address at our own collector and used the web interface as a person would. We gave it the task, clicked the thumbs-up, picked the category Task result and pressed Submit. A 90 KB upload arrived with the whole session in it. Closing the dialog with the cross sent nothing. Taking back an earlier rating did send a record.

Normal use sent nothing: 0.1.5 made no outside connection during ten minutes of use or during the headless task. It did write a permanent anonymous ID to its home folder on first launch, which 0.1.1 did not. With DSH_TELEMETRY_DISABLED=1 the same rating sent nothing, and the ID file was not written. That switch appears in the package files but not in the main README, where the word telemetry does not appear at all.

Case study

One thumbs-up, 90 KB


The feedback dialog gives the only warning, as grey placeholder text in the notes box: "Your submission will include the current conversation log." It disappears when the user starts typing. It does not say where the log goes or what is in it. The plugin list in Settings shows the telemetry plugin as Enabled whether or not it will upload anything.



The feedback dialog, 21 Sep 2026. The only warning is the grey placeholder in the notes box.

FIGURE 10 What the 90 KB upload was made of

Blue is the header that carries the tool definitions, yellow is the conversation itself, grey is session housekeeping. The 24 records in one feedback upload, by size. Half of it is the header that carries every tool definition. The rest is the conversation: the system prompt, the prompt, the replies, and the file the agent wrote, with its contents and its path.



Source: captured payload, StrideNote test, 21 September 2026. Total 89.8 KB of 90.2 KB decoded; StrideNote, 2026 API keys and the computer name were not in the upload.

1
upload of the whole session after one rating, 0.1.5 default

0
uploads with DSH_TELEMETRY_DISABLED=1, same steps

Nothing in this case study reached DeepSeek. The upload address was pointed at our own collector for every run.

FINDINGS Goose | tested 22 Sep 2026, 1.51.0 | grade A, B

Goose asks before it shares, with Yes selected

The only agent of the sixteen that puts the question first, and the only one whose default answer is the one that sends.

Tested:	22 Sep 2026, 1.51.0	Evidence:	A consent and off switch; B payload	Destination:
	us.i.posthog.com	On by default:	Asks first; Yes pre-selected	Off switch:
	Answer No; GOOSE_TELEMETRY_ENABLED; works			

The first Goose session opens on a consent screen that lists what is collected: operating system and version, Goose version and install method, provider and model, tool names and counts, session metrics and error types. It says conversations, code and tool arguments are never collected. The selector starts on Yes, so pressing Enter agrees.

After Yes, Goose contacted `us.i.posthog.com` during the task. That connection refused our certificate, so we did not read the payload. The property names beside the address in the binary match the consent screen. After No, the whole session made no outside request of any kind. Goose made no update or catalogue requests in either case.

```
Help improve goose
Would you like to help improve goose by sharing anonymous usage data?
This helps us understand how goose is used and identify areas for improvement.

What we collect:
• Operating system, version, and architecture
• goose version and install method
• Provider and model used
• Extensions and tool usage counts (names only)
• Session metrics (duration, interaction count, token usage)
• Error types (e.g., "rate_limit", "auth" - no details)

We never collect your conversations, code, tool arguments, error messages,
or any personal data. You can change this anytime with 'goose configure'.

+ Share anonymous usage data to help improve goose?
  Yes / No
```

The first-run consent screen, 22 Sep 2026. The selector starts on Yes.

FINDINGS Aider | tested 22 Sep 2026, 0.86.2 | grade A

Aider asks one install in ten

Analytics stay off for nine installs in ten and the tenth is asked, with the sampling decided on the machine from a random identifier.

Tested:	22 Sep 2026, 0.86.2	Evidence:	A Opted in, captured at our collector	Destination:
	us.i.posthog.com, if you agree	On by default:	No	Off switch: --analytics-disable; works

Aider writes a random ID on first launch and uses it to decide whether to ask for analytics. Only IDs that fall in a ten per cent sample see the question (Aider-AI, 2026). Everyone else is never asked and sends nothing. We opted in on purpose and moved the address to our collector. The events counted launches, the number of files in the repository, token counts and cost, with the operating system and Python version. Aider replaced the local model's name with REDACTED and turned off location lookup.

Without analytics, Aider checked PyPI for a new version and downloaded LiteLLM's model price list from GitHub, which has no switch. The repository's last change was on 22 May 2026.

Agents that borrow, fetch or install

Five agents reported no usage but reached out for other reasons: to update themselves, to fetch catalogues and tools, or in one case to use another program's sign-in.

FINDINGS Hermes Agent | tested 21 Sep 2026, 0.21.1 | grade A

Hermes Agent refreshes Claude Code's sign-in at launch

It reports no usage of its own, then renews Claude Code's saved sign-in at launch, with no setting that stops it.

Tested:	21 Sep 2026, 0.21.1	Evidence:	A Capture in hold mode	Destination:	
	platform.claude.com, plus five update and catalogue hosts	On by default:	Yes	Off switch:	
	None found for the sign-in refresh				

Hermes keeps its own telemetry off by default. Its shared metrics need two separate opt-ins and a recorded consent window (Nous Research, 2026a). On first launch it still made nine requests before we typed anything. Six were ordinary: two for an update check against GitHub, two for copies of its model catalogue, one for the models.dev registry and one for OpenCode's free model list. Two tried to download tirth, a third-party security scanner, and its checksum file from a GitHub account outside the project.

The other was a sign-in refresh to Anthropic. Hermes read the refresh token that Claude Code stores in the macOS Keychain or in ~/.claude, and posted it to platform.claude.com. Our test home configured only LM Studio. The request used the user agent axios/1.7.9. A comment in Hermes's source explains why: Anthropic rejects token requests whose user agent starts with claude-code, and "the real CLI uses bare axios there" (Nous Research, 2026b).

Hermes documents this as a feature, "Auto-detect Claude Code credentials (if you already use Claude Code)" (Nous Research, 2026c). The documentation does not say it happens at every first launch when Anthropic is not in use. We turned off every startup setting Hermes documents and the refresh still ran. Our proxy held the request, so Anthropic never received it, and the captured token was removed from our records.

OpenCode updated itself six seconds after launch

No analytics and no crash reporter in the command line, and a new version installed over the old one before anything was typed.

Tested:	21 Sep 2026, 1.18.30 then 1.18.31	Evidence:	A Capture in hold mode	Destination:
	registry.npmjs.org, models.opencode.ai	On by default:	Yes	Off switch:
	OPENCODE_DISABLE_AUTOUPDATE, OPENCODE_DISABLE_MODELS_FETCH; both work			

OpenCode's command line has no usage analytics and no crash reporter. The Sentry address that edition 1 found belongs to the desktop app. At launch it checks npm for a new version and fetches its model catalogue from models.opencode.ai. Both requests carry the version and tracing headers, and neither carries a body. The documentation says the catalogue comes from models.dev.

On its first launch in our test, OpenCode installed version 1.18.31 over 1.18.30 within six seconds, before anything was typed. Its configuration page says it "will automatically download any new updates when it starts up" (OpenCode, 2026a). With the two documented switches set, OpenCode made no outside request at all and the task still completed.

OpenClaw downloads a helper program past the proxy

No usage reporting, and a helper program fetched from GitHub on a connection that ignored the proxy we had set for it.

Tested:	22 Sep 2026, 2026.9.5	Evidence:	A task and setup; B download	Destination:
	api.github.com, github.com	On by default:	First launch only	Off switch: None found

Setup and the task sent nothing but model traffic. On the first chat launch in a new home folder, OpenClaw connected directly to api.github.com and github.com and saved an executable called fd, a file-search tool, into its agent folder. The download ignored the proxy settings, including Node.js's own option to honour them, so we saw the connections and the file but not the requests. It needs Node.js 24.16 or newer. It also writes a log to a folder in /tmp, outside its home folder, with the computer's name and full paths in it.

OpenHands pulls a repository on every launch

Nothing about the user, and a repository cloned on every launch, so the machine changes each time the agent starts.

Tested: 22 Sep 2026, 1.16.0 **Evidence:** **A** startup requests; **B** clone **Destination:** github.com, pypi.org, raw.githubusercontent.com **On by default:** Yes **Off switch:** None found for the clone or the version check

The OpenHands command line runs the agent on the machine without Docker. It keeps a local copy of the public skills repository, OpenHands/extensions, cloning it on first run and pulling it on every run after. The command line turns this on in its code (OpenHands, 2026). It also checks PyPI for a new version and downloads LiteLLM's price list. The analytics exporter in its SDK is for hosted deployments, requires consent, and was not active.

Grok Build checks for updates and a changelog

An update check and a changelog, nothing about the session, and a documented switch for each.

Tested: 22 Sep 2026, 1.0.40 **Evidence:** **A** Capture in hold mode **Destination:** x.ai, storage.googleapis.com **On by default:** Yes **Off switch:** GROK_DISABLE_AUTOUPDATER for the update check; none for the changelog

Grok Build ran with a local model and no xAI account. At launch it checked x.ai for the current version and downloaded this version's changelog. Its binary contains product analytics and a session-trace upload, and its documentation describes switches for both (xAI, 2026a) and for the update check (xAI, 2026c). Neither fired in our tests. The official installer is a script piped to the shell that edits the shell profile. We read it and downloaded the binary it points to instead.

Agents that sent nothing of their own

Three agents sent nothing beyond model traffic, or only a price list with a working switch.


ZeroClaw, smolagents and Google ADK

Three agents that sent nothing of their own: the only request came from a library fetching a price list, and one setting stopped it.

ZeroClaw made no outside request during setup, idle or the task, and asks before it writes a file. smolagents, a Python library, sent nothing beyond its calls to the model. Google ADK sent nothing itself. The LiteLLM library it uses to reach a local model downloaded its model price list from GitHub when loaded, and LITELLM_LOCAL_MODEL_COST_MAP=True stopped that. With a hosted model, ADK and smolagents would send prompts to that model's provider, which is what they are for.


FIGURE 11 Requests before you type anything

Every outside request an agent made from launch until it was ready for input, at its default settings, with no account and a local model. Sorted lowest to highest. DeepSeek Harness, ZeroClaw and smolagents made none. Kilo Code's five refusals are Homebrew's; OpenClaw's two went to GitHub directly, past the proxy.



Source: first-run capture files in runs/, StrideNote tests, 21 and 22 September 2026. OpenJarvis is not shown: edition 1 did not measure launch requests.

StrideNote, 2026



5


Seven leads looked like telemetry and were not

Seven things that looked like telemetry in the capture and turned out not to be.

Each of these would have made a stronger headline than the finding it turned into.

Lead	What it turned out to be
OpenCode contacts api.opencode.ai	An address match on shared Cloudflare space. The proxy named the host as models.opencode.ai, the model catalogue.
Codex CLI ships Sentry	A crash-reporting address is in the binary. No Sentry request appeared in any run. We did not cause a crash.
Kilo Code telemetryEnabled	The hits belonged to a bundled TypeScript language server's own typings installer, not to Kilo's analytics.
Pi mentions sentry and analytics	Provider tables and names of extension hooks. Pi has no crash reporter or analytics library.
DeepSeek Harness 0.1.1 telemetry Enabled	The plugin list shows the plugin as Enabled because it is mounted. Its mode in 0.1.1 is DISABLED, and it sent nothing.
Hermes ships a telemetry endpoint	It does, at telemetry.nousresearch.com. It needs two opt-ins and a consent window, and both default to off.
Grok Build ships Mixpanel and trace upload	Both are in the binary. Neither fired with a local model and no account.

The rule these point to is the one from edition 1. A string shows what could happen. Only a recorded request shows what did.



6


Five things we got wrong, and what we did about them

The mistakes we made during testing, including the run where our own safeguard did not hold.

A report that asks readers to trust its method should show where the method slipped.

- One run left the proxy out. While checking the Cline command line's settings screen, we started it without the capture proxy. It opened three real connections, to `api.cline.bot`, `otel.cline.bot` and one Cloudflare address, and sent its launch events. No task was given, so no prompt or file content went. The hardware-derived device ID did.
- An install landed in the wrong folder. Installing DeepSeek Harness 0.1.5 into an empty test folder, npm found a package file in the home folder and installed there instead. We removed the 501 packages it added. Three shared libraries stayed on newer compatible versions because the old lock file was overwritten.
- The test updated the tester's own copy. OpenCode's self-update replaced the command line installed on the test machine, leaving it one version ahead of the desktop app.
- A capture held a live credential. Hermes's sign-in refresh put a real refresh token into our capture log. It was never forwarded. We redacted it, deleted the raw file and changed the capture code to redact tokens and passwords before writing.
- A collector dropped the first upload. Our first DeepSeek Harness collector could not read the chunked body of the first upload. We fixed it and repeated the steps on a new session.

We also re-signed our copy of the Cline command line so that macOS would run it, after its shipped signature failed validation in two releases. That changed only our test copy.




7

The switches that worked, and what they left behind

Every switch we found, whether it worked, and what each agent still did with it set.

FIGURE 12 What the documented settings achieved, agent by agent

Requests at launch at default settings (pale dot) and with every switch the maker documents set (solid dot). Seven agents reach zero. Five still send something. Two have no switch. Sorted by what they sent at default, lowest first.



Source: capture files in runs/, StrideNote tests, 21 and 22 September 2026. DeepSeek Harness counts uploads after one rating. OpenJarvis was not assessed.

StrideNote, 2026


Every setting here was tested on the version in Table 1, except OpenHands's LiteLLM setting, which we verified on Google ADK and did not rerun on OpenHands. Left over means requests that still went out with every documented switch set.

Seven agents can be made fully quiet with the settings in the table below: OpenCode, Pi, DeepSeek Harness, Goose, Google ADK, ZeroClaw and smolagents. Eight cannot, because at least one request has no switch that we found. OpenJarvis was not assessed for this.

A setting that works today can change in the next release. DeepSeek Harness changed its default between two releases a month apart. The checks in edition 1's Section 7 still apply: watch the sockets at launch, read the settings file, and look for the address a reporter sends to.

Agent	Settings we used	Left over
OpenCode	OPENCODE_DISABLE_AUTOUPDATE=1, OPENCODE_DISABLE_MODELS_FETCH=1	Nothing

Agent	Settings we used	Left over
Pi	PI_OFFLINE=1 or --offline	Nothing
DeepSeek Harness	DSH_TELEMETRY_DISABLED=1	Nothing
Goose	Answer No at first run, or GOOSE_TELEMETRY_ENABLED: false	Nothing
Google ADK	LITELLM_LOCAL_MODEL_COST_MAP=True	Nothing
ZeroClaw, smolagents	None needed	Nothing
Codex CLI	otel.metrics_exporter = "none", check_for_update_on_startup = false , analytics.enabled = false , feedback.enabled = false , features.plugins = false , features.remote_plugin = false	The announcement file from GitHub
Kilo Code	KILO_TELEMETRY_LEVEL=off , KILO_DISABLE_MODELS_FETCH=1 , "autoupdate": false	Four requests to api.kilo.ai: two for the model list, two for model defaults
Aider	--analytics-disable , --no-check-update	LiteLLM's price list
Grok Build	GROK_DISABLE_AUTOUPDATER=1	The changelog from x.ai
Hermes Agent	updates.check: false , model_catalog.enabled: false , security.tirith_enabled: false	Claude Code sign-in refresh, OpenCode model list, models.dev (first launch)
OpenHands	LITELLM_LOCAL_MODEL_COST_MAP=True	Skills repository pull, PyPI version check (not rerun)
OpenClaw	None found	The fd download on first launch
Cline (command line)	None found; DO_NOT_TRACK=1 had no effect	All telemetry events and the feature-flag request
OpenJarvis	[analytics] enabled = false (edition 1)	Analytics stopped; other traffic not assessed




8

What these findings do not show

What one machine, one model and one task can and cannot show.

The gaps are listed so they are not read as clean results.

- Local model only. With a hosted model, every agent sends the prompt to that model's provider by design. We tested what the agent sends on its own account.
- No accounts. Kilo Code, Cline, Grok Build and others add signed-in data and events. We signed in to nothing, so those paths are code findings at most.
- One machine, one operating system, one version per agent, on one or two days. Behaviour on Windows, Linux and other versions may differ.
- Six flows refused our certificate: Goose's analytics, Codex's metrics transport and one Codex request to github.com, the two Homebrew hosts that Kilo Code's update check reached, and OpenHands's clone. OpenClaw's download did not use the proxy at all. For those we know where the data went, not what it held. For Codex we read the same pipeline at our own collector.
- Command line forms only. The Cline and Kilo Code editor extensions, the OpenCode and Hermes desktop apps and the OpenHands web interface, which needs Docker, were not run.
- Short windows. Our runs lasted two to ten minutes. OpenCode refreshes its catalogue every 60 minutes, and a report on a longer timer would not appear.
- OpenJarvis was not retested. Its result is from July.
- Selection. The sixteen are the agents we chose to test, not a sample of the market. The counts are not rates for agents in general.



9

Six tests that close the gaps


The six tests that would close the gaps this edition leaves open.

Each item closes a gap named in Section 8.

Test	Why	Target grade
Signed-in runs	Kilo Code, Cline and Grok Build define events that only fire with an account	A
Editor extensions	Run Cline and Kilo Code inside the editor and separate the editor's own traffic	A
Desktop apps	OpenCode desktop, whose crash reporter edition 1 found, and the Hermes app	A
A working day per agent	Replace minutes with hours to catch reports on long timers	A
Payloads that refused us	Goose and OpenClaw, with a capture method that does not change system trust	A
Release tracking	Re-run the fixed task on each new release of the sixteen and record any change in defaults	As 2026

FIGURE 13 Where the next edition's tests fall

The six tests in the table above, and the agents each one applies to. Two apply to all sixteen. The other four close gaps that belong to a handful of agents.



Source: Sections 8 and 9 of this report.

StrideNote, 2026

The dataset

The same rows are published as `agent-telemetry-2026.csv`. Each row is one finding, not one agent. The videos, screen captures and request logs behind each row are kept with the dataset.

Agent	Version	Date	Gr.	Destination	Trigger	What is sent	Default	Off switch
OpenJarvis	desktop-v1.0.2	2026-07-08	A	PostHog host 34.231.106.201.sslip.io, hard-coded project key	During use, 30 second flush	Event chat_session_ended: turn count, tokens in and out, latency p50 and p95, tool counts, error count, duration, hashed model name, OS, Python ...	yes	[analytics] enabled = false in ~/.openjarvis/config.toml
OpenJarvis installer	install.sh	2026-07-08	C	Analytics endpoint	During install	Install progress	yes	None; OPENJARVIS_ANALYTICS_HOST redirects the request
OpenCode CLI	1.18.31	2026-09-21	A	registry.npmjs.org /opencode-ai/latest	At launch	GET with empty body; User-Agent opencode/<version>; per-request traceparent and b3 trace IDs; IP address	yes	autoupdate false or notify in config; OPENCODE_DISABLE_AUTOUPDATE
OpenCode CLI	1.18.31	2026-09-21	A	registry.npmjs.org (package download)	At launch when a newer version exists	Package download; installed 1.18.31 over 1.18.30 within 6 s, no prompt	yes	autoupdate false in config; OPENCODE_DISABLE_AUTOUPDATE
OpenCode CLI	1.18.31	2026-09-21	A	models.opencode.ai /api.json	At launch when the cache is stale, then every 60 minutes	GET with empty body; User-Agent opencode/latest/<version>/cli; per-request trace IDs; IP address	yes	OPENCODE_DISABLE_MODELS_FETCH
DeepSeek Harness	0.1.5-rc.2	2026-09-21	A	harness-telemetry.deepseek.com /v1/logs (shipped default; ...)	Submitting a thumbs rating (category only is enough), text ...	Full session log: system prompt, tool schemas, user prompts, tool calls and results including file contents, assistant replies, absolute workspace ...	yes (FEE DBA CK_ONLY)	DSH_TELEMETRY_DISABLED (any value) or DSH_TELEMETRY_MODE=DISABLED
DeepSeek Harness	0.1.5-rc.2	2026-09-21	A	\$DSH_HOME/.anonymous-user-id (local)	First launch in default mode, before any feedback	Random UUID stored locally; attached as user.id to every upload	yes	DSH_TELEMETRY_DISABLED prevents the file
DeepSeek Harness	0.1.1-rc.2	2026-09-21	B	None observed; telemetry plugin shipped with mode DISABLED	n/a	Nothing observed	no	n/a

Agent	Version	Date	Gr.	Destination	Trigger	What is sent	Default	Off switch
Pi	0.87.0	2026-09-21	A	pi.dev /api/report-install	First launch of each new version	GET with version in the query; User-Agent pi/<version> (OS; node version; CPU arch); IP address	yes	PI_TELEMETRY=0 or enableInstallTelemetry false; PI_OFFLINE=1
Pi	0.87.0	2026-09-21	A	pi.dev /api/latest-version	Every launch	GET, no body; same User-Agent; IP address	yes	PI_OFFLINE=1 or --offline
Pi	0.87.0	2026-09-21	A	github.com /sharkdp/fd/releases/latest	Launch when the fd helper binary is missing	GET, no body; User-Agent pi-coding-agent	yes	PI_OFFLINE=1
Hermes Agent	0.21.1	2026-09-21	A	platform.claude.com /v1/oauth/token	First launch of a Hermes home, with Anthropic not configured	POST grant_type=refresh_token with Claude Code's stored refresh token (read from macOS Keychain or ~/.claude/.credentials.json) and client_id; ...	yes	None found
Hermes Agent	0.21.1	2026-09-21	A	github.com NousResearch/hermes-agent.git (git ls-remote)	Launch, cached 6 hours	Git smart-HTTP ref listing; User-Agent git/<version>; IP address	yes	updates.check: false
Hermes Agent	0.21.1	2026-09-21	A	hermes-agent.nousresearch.com /docs/api/model-catalog.json; ...	Launch; cache 20 minutes	GET, no body; User-Agent hermes-cli/<version>	yes	model_catalog.enabled: false
Hermes Agent	0.21.1	2026-09-21	A	models.dev /api.json	Launch	GET, no body; User-Agent python-requests	yes	None (url override only)
Hermes Agent	0.21.1	2026-09-21	A	github.com sheeki03/tirth/releases (binary + checksums)	First launch when the tirth scanner is missing	GET download of a third-party security-scanner binary	yes	security.tirth_enabled: false
Hermes Agent	0.21.1	2026-09-21	A	opencode.ai /zen/v1/models	First launch	GET, no body; User-Agent HermesAgent/<version>	yes	None found
Codex CLI	0.155.1	2026-09-22	A	ab.chatgpt.com /otlp/v1/metrics (Statsig)	During use, batched; also at idle	52 usage metric series: API/tool/turn counts and durations, token counts by type, model name (even local), sandbox policy, is_git, terminal, feature ...	yes	otel.metrics_exporter = "none"
Codex CLI	0.155.1	2026-09-22	A	api.github.com /repos/openai/codex/releases/latest	Launch	GET, no body; User-Agent with exact macOS version, arch and terminal	yes	check_for_update_on_startup = false
Codex CLI	0.155.1	2026-09-22	A	chatgpt.com /backend-api/plugins/featured and ...	Launch	GET, no body; User-Agent with OS version	yes	features.plugins = false and features.remote_plugin = false
Codex CLI	0.155.1	2026-09-22	A	raw.githubusercontent.com /openai/codex/main/announcement_tip.toml	Launch	GET, no body	yes	None found

Agent	Version	Date	Gr.	Destination	Trigger	What is sent	Default	Off switch
Aider	0.86.2	2026-09-22	A	us.i.posthog.com (PostHog)	Only if the user accepts the opt-in prompt, which is shown ...	Events launched, repo (file count), auto_commits, cli session, message_send (token counts, cost); OS and version, Python version, arch, Aider ...	no	--analytics-disable
Aider	0.86.2	2026-09-22	A	pypi.org /pypi/aider-chat/json	Launch	GET, no body; User-Agent python-requests	yes	--no-check-update
Aider	0.86.2	2026-09-22	A	raw.githubusercontent.com ...	Launch, cached 24 hours	GET, no body	yes	None found
OpenClaw	2026.9.5	2026-09-22	B	api.github.com and github.com (fd release download)	First launch of a new home	Release lookup and download of the fd 10.5.0 arm64 binary into the agent folder; ignores HTTPS_PROXY and NODE_USE_ENV_PROXY	yes	None found
OpenClaw	2026.9.5	2026-09-22	A	None (LM Studio only)	Onboarding, launch, task	Nothing observed	n/a	n/a
Goose	1.51.0	2026-09-22	B	us.i.posthog.com /capture/ (PostHog)	During sessions, after the user accepts the first-run ...	Per disclosure and binary: OS, arch, version, install method, provider and model, extension and tool names with counts, session number and metrics, ...	asks; default answer Yes	Answer No at the prompt; GOOSE_TELEMETRY_ENABLED false; GOOSE_TELEMETRY_OFF
smolagents	1.26.0	2026-09-22	A	None (LM Studio only)	Task run	Nothing observed	n/a	n/a
Google ADK (via LiteLLM)	2.9.2 / litellm 1.102.0	2026-09-22	A	raw.githubusercontent.com ...	When LiteLLM is imported	GET, no body	yes	LITELLM_LOCAL_MODEL_COST_MAP=True
Kilo Code CLI	7.7.7	2026-09-22	A	us.i.posthog.com /batch/ (PostHog, GeoIP on)	Launch (CLI Start); further session events are defined for ...	CLI Start: app name and version, platform, OS and exact OS version, arch, random distinct_id from ~/.local/share/kilo/telemetry -id; IP address ...	yes	KILO_TELEMETRY_LEVEL set to anything other than all
Kilo Code CLI	7.7.7	2026-09-22	A	formulae.brew.sh and analytics.brew.sh via Homebrew; ...	Launch (install-method and update check)	Kilo spawns npm list and Homebrew; Homebrew's own requests, including its analytics endpoint	yes	autoupdate false in kilo.json
Kilo Code CLI	7.7.7	2026-09-22	A	api.kilo.ai /api/openrouter/models and /api/defaults	Launch	GET, no body	yes	None found
ZeroClaw	0.8.5	2026-09-22	A	None (LM Studio only)	Quickstart, launch, task	Nothing observed	n/a	n/a
Grok Build	1.0.40	2026-09-22	A	x.ai /cli/stable (fallback ...	Launch	GET, no body	yes	GROK_DISABLE_AUTOUPDATER=1

Agent	Version	Date	Gr.	Destination	Trigger	What is sent	Default	Off switch
Grok Build	1.0.40	2026-09-22	A	x.ai /cli/changelogs/<version>.external.md and .json	Launch	GET, no body	yes	None found
Grok Build	1.0.40	2026-09-22	A	Mixpanel and SpaceXAI events or trace upload (in binary; not observed)	n/a with a local model and no login	Nothing observed	not observed	[features] telemetry false; GROK_TELEMETRY_ENABLED; trace_upload false; ...
OpenHands CLI	1.16.0 (SDK 1.21.0)	2026-09-22	B	github.com OpenHands/extensions (git clone, then pull)	Every launch (clone on first run, pull after)	Git clone or fetch of the public skills registry	yes	None found (load_public_skills=True is hard-coded in the CLI)
OpenHands CLI	1.16.0 (SDK 1.21.0)	2026-09-22	A	pypi.org /pypi/openhands/json	Launch (splash screen)	GET, no body	yes	None found
OpenHands CLI	1.16.0 (SDK 1.21.0)	2026-09-22	A	raw.githubusercontent.com ...	Launch (LiteLLM import)	GET, no body	yes	LITELLM_LOCAL_MODEL_COST_MAP=True
Cline CLI	3.0.64	2026-09-22	A	otel.cline.bot/v1/logs (OTLP JSON)	Launch and throughout every task	37 lifecycle, task and agent events (24 kinds) in one short task: provider, model name, token counts, tool names, mode, has_git, OS and exact OS ...	yes	None found for the CLI (DO_NOT_TRACK ignored); extension: Cline Telemetry ...
Cline CLI	3.0.64	2026-09-22	A	data.cline.bot/flags/ (PostHog feature flags)	Launch	POST with PostHog key, distinct_id = the hardware-derived device ID, flag keys; geoip_disable true	yes	None found

Sources

StrideNote

StrideNote (2026a) Local AI Telemetry Report 2026. Edited by StrideNote Studio. StrideNote Telemetry Report, edition 1. Zenodo. <https://doi.org/10.5281/zenodo.22848632>

StrideNote (2026b) OpenCode contacts Sentry before you type anything, and two other findings that were not real. 8 August. <https://stridenote.net/opencode-telemetry/>

Primary documents

AAIF (2026) goose v1.51.0 release. <https://github.com/aaif-goose/goose/releases/tag/v1.51.0>

Aider-AI (2026) aider/analytics.py. <https://github.com/Aider-AI/aider/blob/main/aider/analytics.py>

BerriAI (2026) LiteLLM. <https://github.com/BerriAI/litellm>

Cline (2026a) Cline Telemetry. <https://docs.cline.bot/more-info/telemetry>

DeepSeek (2026a) @deepseek-ai/dsh-session-telemetry-otel 0.1.5-rc.2, README. <https://www.npmjs.com/package/@deepseek-ai/dsh-session-telemetry-otel>

DeepSeek (2026b) @deepseek-ai/dsh-session-telemetry 0.1.5-rc.2, README. <https://www.npmjs.com/package/@deepseek-ai/dsh-session-telemetry>

Earendil (2026) Pi coding agent 0.87.0, README (Telemetry and update checks) and docs/settings.md. <https://github.com/earendil-works/pi>

Google (2026) Agent Development Kit for Python. <https://github.com/google/adk-python>

Hugging Face (2026) smolagents. <https://github.com/huggingface/smolagents>

Kilo Code (2026) Settings. <https://kilo.ai/docs/getting-started/settings>

Nous Research (2026a) hermes_cli/config_defaults.py. https://github.com/NousResearch/hermes-agent/blob/main/hermes_cli/config_defaults.py

Nous Research (2026b) agent/anthropic_credentials.py. https://github.com/NousResearch/hermes-agent/blob/main/agent/anthropic_credentials.py

Nous Research (2026c) Providers. <https://github.com/NousResearch/hermes-agent/blob/main/website/docs/integrations/providers.md>

OpenAI (2026a) Codex configuration reference. <https://developers.openai.com/codex/configuration-reference>

OpenClaw (2026) openclaw. <https://github.com/openclaw/openclaw>

OpenCode (2026a) Config. <https://opencode.ai/docs/config/>

OpenCode (2026b) CLI. <https://opencode.ai/docs/cli/>

OpenHands (2026) OpenHands CLI. <https://github.com/OpenHands/OpenHands-CLI>

xAI (2026a) Grok Build user guide: configuration. <https://github.com/xai-org/grok-build/blob/main/crates/codegen/xai-grok-pager/docs/user-guide/05-configuration.md>

xAI (2026b) Grok Build user guide: monitoring usage. <https://github.com/xai-org/grok-build/blob/main/crates/codegen/xai-grok-pager/docs/user-guide/24-monitoring-usage.md>


xAI (2026c) Grok Build user guide: config reference. <https://github.com/xai-org/grok-build/blob/main/crates/codegen/xai-grok-pager/docs/user-guide/26-config-reference.md>

zeroclaw-labs (2026) ZeroClaw v0.8.5 release. <https://github.com/zeroclaw-labs/zeroclaw/releases/tag/v0.8.5>

Appendix C

Glossary

Agent	A program that uses a language model to act on your machine: read and write files, run commands, call tools.
Capture proxy	A program on the same machine that sits between the agent and the internet and records each request.
Hold mode	Our proxy setting that records each request and answers it on the machine, so nothing is passed on.
Collector	A small server on our machine that receives telemetry when an agent lets us change its reporting address.
OpenTelemetry (OTLP)	An open standard for sending logs, metrics and traces. Several agents use it to report usage.
PostHog, Statsig	Product analytics services. Agents send them events with an ID so usage can be counted.
Hardware UUID	A fixed identifier for a Mac, set at the factory. A hash of it stays the same across reinstalls.
Refresh token	A long-lived credential an app exchanges for short-lived access tokens without asking the user to sign in again.
User agent	A line each request carries that names the software sending it, and often its version and operating system.
Marker word	STRIDEMARK-4417, placed in the task so that any request carrying the prompt could be found by search.



Where the data went.

StrideNote.net is an independent research publication on AI and modern technological systems. We follow the stories, run the tools, read the code and publish the evidence.

Stridenalysis is our research desk. Its investigations, reports and insights are free to read, and every report is published with the data behind it.

Local Agent Telemetry Report 2026, StrideNote Telemetry Report, edition 2. Published by StrideNote.net and edited by StrideNote Studio. This report describes software behaviour observed on one machine during specific test windows in July and September 2026. It is provided for information only and is not security, legal or investment advice. The report and its dataset are licensed CC BY 4.0. Product names belong to their owners, and no company named here reviewed or paid for this report. Set in Geist and Inter.

READ

stridenote.net
stridenote.substack.com

FOLLOW

x.com/stridenote
linkedin.com/company/stridenote
bsky.app/profile/stridenote.net
youtube.com/@stridenote
huggingface.co/stridenote

CONTACT

hello@stridenote.net



CC BY 4.0

Edition 2 | September 2026
DOI 10.5281/zenodo.22969237