

# Repo Triage at Scale and New OpenClaw Control Releases

Coding Agents Alpha Tracker

2026-04-26

## Repo Triage at Scale and New OpenClaw Control Releases

*By Coding Agents Alpha Tracker • April 26, 2026*

The clearest signal today is practical: steipete turned repo maintenance into a staged agent pipeline — mass close first, cluster later. Also in this brief: acpx 0.6.0, CodexBar 0.23, pragmatic model-routing notes, remote agent execution, and a DeepSeek V4 clip worth a minute.

### TOP SIGNAL

Repo maintenance produced the day's most usable agent pattern. Peter Steinberger says `clawsweeper` runs **50 Codex instances in parallel 24/7** to deeply scan issues and PRs and close items that are already implemented or make no sense, and that pass closed around **4,000 issues** in a day; he then describes **intent-based clustering** as the next pass, with **Project Clowfish** running alongside **Terminator Claw** to group relevant issues and PRs [1, 2, 3].

If you manage a noisy repo, the copyable takeaway is simple: **pre-clean first, cluster second** [2, 3].

### TOOLS & MODELS

- **acpx 0.6.0** — agent control layer for Codex and Claude. New bits: Claude system-prompt controls, session pruning, embeddable turn handles, `--no-terminal`, persistent-session fixes, WSL cwd translation, queue hardening, and clearer error hints. Release notes [4]
- **CodexBar 0.23** — adds Mistral support, Claude Designs/Daily Routines usage, Cursor Extra usage, GPT-5.5 pricing, cleaner widgets and menus, and reliability fixes. Release notes [5]

- **Current practitioner model split:** McKay Wrigley says **GPT 5.5** is the single model he'd keep for engineering, **Codex** has been the gold-standard coding app since 5.3, **Claude Code** is still T1, and **Claude** remains stronger for his non-coding agent work; he also says the updated Claude Code app was good enough to move him out of the terminal [6].
- **Packaging matters too:** Riley Brown's take is that Anthropic's **Agent SDK** is probably under-adopted partly because the name hides what developers would use it for [7].
- **Cost check:** Theo says a single **Opus** prompt can burn more than an **\$8** chat tier, which fits swyx's "code is not cheap" framing around Matt Z Carey's new agent talk [8, 9, 10].

## WORKFLOWS & TRICKS

- **Copy steipete's repo-triage loop:** (1) run a pre-clean pass that closes fixed or nonsense issues with parallel agents, (2) then run intent-based clustering on what survives, with Project Clowfish and Terminator Claw grouping related issues and PRs [1, 2, 3].
- **Use the README as the status surface.** Steinberger says **clawsweeper** updates README.md as it works instead of maintaining a separate dashboard [11].  
"Readme is the new dashboard." [11]
- **Keep long-running agents on a remote box.** Theo set up a server he can SSH into so agents keep running reliably even when his local internet is bad. If your sessions keep dying because of connectivity, move execution off your laptop first [12].
- **Route models by task.** McKay's current practical split: Codex/GPT for coding, Claude for non-coding agent tasks [6].

## PEOPLE TO WATCH

- **Peter Steinberger** — highest-signal builder in today's notes. He is shipping the control layer (**acpx**), the integration surface (**CodexBar**), and the repo-maintenance workers behind OpenClaw-style triage [4, 5, 1].
- **Theo** — worth following for operator notes, not benchmark chatter: today it was remote agent execution under bad connectivity and the cost of premium models relative to cheap chat tiers [12, 8].
- **Matt Z Carey** — new talk drop: Every API is a Tool for Agents. swyx singled it out as another serious "code is not cheap" take, which makes it a useful counterweight to the idea that agentic code generation is cheap by default [10, 9].
- **Riley Brown** — useful when product packaging starts to matter as much as model quality; his Agent SDK naming comment is a good reminder that discoverability changes adoption [7].

## WATCH & LISTEN

- **02:55-03:56** — **DeepSeek V4 in one minute.** Matthew Berman walks through the concrete shipped specs: Pro vs. Flash, **1M token context**, MoE sizes, and **33T-token** training. Worth the minute if you just want the release facts [13].



*My Honest Thoughts about Deepseek (2:55)*

- **04:16-04:58** — **Why this open release matters.** Berman's summary is that DeepSeek V4 beats current open models in math, STEM, and coding, sits just behind the very top closed models on benchmarks like SweepBench Verified, and does it at a fraction of the price [13].



*My Honest Thoughts about Deepseek (4:16)*

## PROJECTS & REPOS

- **clawsweeper** — steipete’s repo-cleanup worker. It runs **50 Codex** in parallel and reportedly closed around **4,000 issues** in a day, with more in the pipeline [1].
- **acpx v0.6.0** — control tooling for Codex/Claude via agents; this release tightens session management, terminal behavior, and queue reliability [4].
- **CodexBar v0.23** — coding-AI integration tool with new Mistral support, GPT-5.5 pricing, and reliability fixes [5].
- **Project Clowfish + Terminator Claw** — names to watch inside the same repo-triage stack; Steinberger and Vincent Koc describe them as the clustering layer that finds related issues and PRs after the initial cleanup pass [2, 3].

*Editorial take: today’s biggest edge was operational, not magical — run agents somewhere stable, route the right model to the right job, and turn messy repo work into staged pipelines instead of one giant prompt. [12, 6, 2, 1]*

---

## Sources

1. X post by @steipete

2. X post by @steipete
3. X post by @vincent\_koc
4. X post by @steipete
5. X post by @steipete
6. X post by @mckaywrigley
7. X post by @rileybrown
8. X post by @theo
9. X post by @swyx
10. X post by @mattzcarey
11. X post by @steipete
12. X post by @theo
13. My Honest Thoughts about Deepseek