

# Cursor Details Its Async-Agent Playbook; DevTools MCP Adds Agent QA Skills

Coding Agents Alpha Tracker

2026-04-13

## Cursor Details Its Async-Agent Playbook; Dev- Tools MCP Adds Agent QA Skills

*By Coding Agents Alpha Tracker • April 13, 2026*

Aman Sanger shared rare production metrics and concrete operating patterns for cloud coding agents, from artifact-first review to planner/subplanner trees. Also in today's signal: DevTools MCP adds built-in web quality checks, Kent C. Dodds shows an agent-to-agent repair loop, and Addy Osmani puts guardrails around parallel-agent overload.

### TOP SIGNAL

Cursor founder Aman Sanger shared rare production datapoints: async coding agents running in cloud VMs are already generating 30% of merged PRs internally, and by the end of 2025 agent requests exceeded tab accepts even though tab can fire on every keystroke [1]. This is not toy output: he cites an eight-hour React-to-Rust video-rendering refactor that made the system 25x faster, plus a 10,000-line PR for network policy controls [1]. The actionable part is the workflow: give long-running agents a full computer, let them test their own work, and review video/report artifacts before you spend time reading code [1].

### TOOLS & MODELS

- **Cursor Cloud Agents: model routing is task routing.** Cursor uses OpenAI models for higher-level planning/orchestration, Gemini and Anthropic for computer use and multimodal work, Anthropic for stronger UI generation, and faster models for simpler subagent tasks where they match performance at lower latency [1].
- **Chrome DevTools MCP just got more useful for agent QA.** Addy Osmani says it now includes Lighthouse performance checks, a memory

leak detection skill, an accessibility debugging skill, an LCP optimization skill, and an experimental CLI [2].

- **Codex UX signal from a power user.** Riley Brown likes Codex because it stays bare-bones, keeps project work in left-sidebar folders with multiple threads, and should auto-spin new threads when a request splits into multiple tasks instead of adding extra interface layers like a Codex Cowork middle step [3].
- **Local model reality check.** Armin Ronacher says local models are not universally there yet, but Qwen3-Coder-Next on a beefy Mac is useful for its speed and decent tool calling; his broader point is not to dismiss local models outright, especially for tasks like STT [4].

## WORKFLOWS & TRICKS

- **Artifact-first review loop for long-running agents.**
  1. Run the agent in a cloud VM with the same kind of tools you'd use locally.
  2. Let it implement and test the change itself.
  3. Review its artifacts first, especially a demo video or research report.
  4. If something looks off, reprompt from the artifact instead of diving straight into the diff.
  5. Only review code once the behavior looks right [1].
- **Break long jobs into a planner tree.** Cursor's working pattern is planner -> subplanners -> workers. Keep each leaf task simple enough to stay within training distribution, cap the outer agent to hundreds of thousands of tokens rather than tens of millions, and use faster models for simpler subagent tasks when they achieve the same result [1].
- **Use event-driven agents, not just ad hoc prompts.** Cursor's automations fire on issues, pages, training runs, PRs, and security events. Their examples: an agent investigates a middle-of-the-night page and proposes a single-click fix, watches training metrics and logs to catch failures early, and finds PR vulnerabilities before shipping [1].
- **Let one agent escalate to another agent.** Kent C. Dodds describes a clean loop with Kody: when Kody hits a problem, he tells it to kick off a Cursor Cloud Agent, and Kody already knows the problem context well enough to compose the repair prompt itself [5, 6].
- **Parallel agents are not free throughput.** Addy Osmani's advice: time-box long agentic sessions like deep work, tighten the scope for each agent, and treat finding your personal ceiling as a skill. Addy's longer note: Your parallel Agent limit [7, 8].

I can fire up four agents in parallel and have them work on four different problems, and by 11am I am wiped out for the day. [9]

- **On unfamiliar platforms, shrink the blast radius.** Trash Dev says he shipped his first iPhone app with 100% AI-generated code despite no prior Swift or iOS experience, but kept it local-only, manually checked

local storage and external calls, and used a separate design tool when the UI was the bottleneck [10].

## PEOPLE TO WATCH

- **Aman Sanger.** Shared production metrics and concrete architecture details: PR share, artifact review, multi-agent decomposition, model routing, and event-driven automations [1].
- **Addy Osmani.** Posted on two useful fronts today: human-in-the-loop constraints for parallel agents and practical agent QA tooling via DevTools MCP [7, 2, 11].
- **Kent C. Dodds.** Showed an agent-to-agent repair loop where the agent using the software can hand off to a cloud repair agent without re-explaining the bug [5, 6].
- **Armin Ronacher.** Shared a grounded local-model take: one concrete setup that works on strong Mac hardware today instead of blanket yes or no claims about local coding models [4].

## WATCH & LISTEN

- **3:48-4:48 — Why async agents need their own computers.** Sanger lays out why long-running agents need cloud VMs with full desktop access and testing, not just editor-side autocomplete [1].



*Building Towards Self-Driving Codebases with Long-Running, Asynchronous*

*Agents (3:48)*

- **6:40-8:14 — Review artifacts before code.** Sanger explains why videos and research reports are faster review surfaces than raw diffs for early agent iterations [1].



*Building Towards Self-Driving Codebases with Long-Running, Asynchronous Agents (6:40)*

- **8:18-9:05 — Zero-Swift iOS app, but still audit the risky path.** Trash Dev explains the 100% AI-coded experiment, then grounds it by keeping the app local-only and manually checking storage and external calls before shipping [10].



*Trash Made a Black Mirror App / The Standup (8:18)*

## PROJECTS & REPOS

- **Chrome DevTools MCP.** New quality-check surface for agents: Lighthouse performance checks, memory leak detection, accessibility debugging, LCP optimization, plus an experimental CLI. Repo: `chrome-devtools-mcp` [2, 11].
- **Cursor's browser prototype.** Sanger says Cursor ran a one-week multi-agent build that spent billions of tokens and tens of thousands of dollars of compute to produce a working, but still far-from-production, browser [1].

*Editorial take: today's serious workflows looked less like autocomplete and more like ops — cloud VMs, event triggers, review artifacts, and explicit limits on human attention. [1, 7]*

---

## Sources

1. Building Towards Self-Driving Codebases with Long-Running, Asynchronous Agents
2. X post by @addyosmani
3. X post by @rileybrown

4. X post by @mitsuhiko
5. X post by @kentcdodds
6. X post by @kentcdodds
7. X post by @addyosmani
8. X post by @addyosmani
9. X post by @lennysan
10. Trash Made a Black Mirror App | The Standup
11. X post by @addyosmani