

Muse Code Makes Cost and Locality Part of the Coding-Agent Stack

Coding Agents Alpha Tracker

2026-08-11

Muse Code Makes Cost and Locality Part of the Coding-Agent Stack

By Coding Agents Alpha Tracker • August 11, 2026

Meta’s Muse Code beta and Muse Glimmer push cheaper terminal and local execution into the coding-agent stack. Practitioner tests and new orchestration tools point to task-level routing, reusable loops, and reviewable autonomy as the practical edge.

TOP SIGNAL

Meta is making cost and locality first-class coding-agent choices. Muse Code beta is a terminal agent for complete software-engineering tasks across large repos—planning changes, writing code, and validating results—powered by Muse Spark 1.2; Riley Brown places it between Opus and GPT-5.6 Tera and calculates roughly \$5.50 in combined input/output cost for Meta’s top model versus \$30 for Opus and \$35 for GPT-5.6. [1]

Muse Glimmer adds a 30B Apache 2.0 agentic model: Simon Willison generated an example with LM Studio’s 18.16GB build and says a 32GB+ machine leaves room for other applications. [2, 3] The practical shift is to measure completed-task cost, latency, and local execution alongside output quality—not simply choose the highest-ranked model. [4, 5, 1]

TRY THIS

- **Route by total task cost, not token price.** In DHH’s follow-on rewrite test, GPT-5.6 Sol High followed Fable’s plan: the first pass was 30% slower, but one follow-up reached parity at \$43; DeepSeek V4 Flash could not get anything working despite many follow-ups. [6] Kimi K3 was stopped after \$60; when DHH let K3 Fast finish, it used \$80 in tokens—about \$55 at standard pricing—and was still slower and more expensive

than Sol High. [7, 8] Start with a fixed acceptance test and a turn/dollar budget for the cheap model, escalate when it stalls, and log total task cost. Riley’s warning is the right accounting rule: low per-token pricing can be erased by extra turns. [1]

- **Make optimization a separate, measured pass.** DHH reports that a GPT Sol optimization run immediately made `ttfx` 63% faster—14× the original `tte`—followed by 16× and then 27× results in later rounds. [9, 10, 11] After the baseline works, give the agent a bounded prompt such as: `Profile this implementation against the original; optimize only measured bottlenecks; run the same benchmark; report before/after and regressions.`
- **Orient before editing a neglected repo.** ThePrimeagen opened a project untouched for more than a year, asked `"yo ai, remind me of ...."`, and got up to speed in about three minutes. [12] Use that first pass to recover architecture, entry points, tests, and unknowns; do not start implementation until the agent can restate the plan.
- **Use a reviewer agent to create a triage map.** Kent C. Dodds asked Devin for a deep review of Kody. The returned review took six minutes and added 96 lines on a roughly 470,000-line codebase, citing zero TODO/FIXME/HACK comments, four uses of `any`, 606 test files, matching security invariants, and a “strong but not perfect” verdict. [13, 14] Treat this as a fast findings inventory for human verification, not as an autonomous approval gate.

WHAT SHIPPED

- **Muse Glimmer open weights.** Meta announced a 30B dense model trained for agentic use cases under Apache 2.0, with a stated 24GB-of-VRAM target; the GGUF build is available on Hugging Face. [2, 4, 15]
- **Cross-harness coordination is productizing.** Spotify launched Xirp, a vendor-neutral environment for managing sessions across Claude, Gemini CLI, and Codex; Spotify says 1,300+ engineers already use it and the service is now available to try. [16] Riley Brown’s Buzz walkthrough shows the complementary team pattern: @mention Codex and Claude Code in one Slack-like thread, create a Cursor agent backed by Kimi K3, add it to channels, and watch which agent is working or requesting approval; the demo agent joined seven channels. [1]
- **Kody v2026.08.10** adds built-in OAuth integrations. Operators can provision shared GitHub, Google, Slack, and similar OAuth apps through `/admin/platform-integrations`; Kent says onboarding feedback made this a necessary shift. [17, 18]
- **loop-library is now open source.** Jason Zhou’s collection packages copyable prompts with loops he says deliver real-world results; he says

most of the listed loops are already running in SuperDesignDev. [19, 20]

- **Durable agent chat got a concrete product treatment.** Addy Osmani highlights Trigger.dev’s new chat agent, which survives refreshes, crashes, and redeploys and can pause for permission before a risky tool action. [21]


GO DEEPER

- **Riley Brown — Meta’s NEW Muse Code is Here and Major Codex Updates:** Watch the install → authentication → terminal `muse` → Wii Bowling flow. The important operational detail is the sandbox tradeoff: the demo configures YOLO mode to remove permission prompts and explicitly gives the agent full control, so reproduce that only in a disposable environment. [1]



Meta’s NEW Muse Code is Here and Major Codex Updates (2:33)

- **LangChain — Create an agent that can browse the web with Managed Deep Agents and Browserbase’s Stagehand:** Study the tool contract—`snapshot`, `run`, and `screenshot`—including the rule to re-snapshot after navigation; then follow the local `mda dev test`, Browserbase session replay, and `mda deploy` flow. The deployment keeps instructions in Context Hub, so prompts can change without redeploying the agent.



[22]

Create an agent that can browse the web with Managed Deep Agents and Browserbase’s Stagehand (3:28)

- **Repo to study** — **ttfx**: Trace the progression from the first measured 63% optimization to 14×, 16×, and 27× claims. It is a compact example of separating agent-generated implementation from repeatable performance passes. [9, 10, 11]

Editorial take: The high-alpha move is not picking one permanent model winner; it is making model routing, context recovery, optimization, review, and permissions explicit stages of one inspectable workflow. [1, 6, 21, 19]

Sources

1. Meta’s NEW Muse Code is Here and Major Codex Updates
2. X post by @alexandr_wang
3. Introducing Muse Glimmer
4. X post by @jack_w_rae
5. X post by @jack_w_rae
6. X post by @dhh
7. X post by @dhh
8. X post by @dhh
9. X post by @dhh
10. X post by @dhh
11. X post by @dhh

12. X post by @ThePrimeagen
13. X post by @kentcdodds
14. X post by @kentcdodds
15. X post by @simonw
16. X post by @SpotifyEng
17. X post by @kodykoala
18. X post by @kentcdodds
19. X post by @jasonzhou1993
20. X post by @jasonzhou1993
21. X post by @addyosmani
22. Create an agent that can browse the web with Managed Deep Agents and Browserbase's Stagehand