

\$750K @ \$12M cap

Delaware C-Corp · June 2026

Context Mode

the substrate AI coding agents run on.

An open-source MCP plugin **290,000+ developers** run locally. A Platform with paid Solutions for engineering organizations. **Same engine, two surfaces.**

Every AI coding agent hits the *same wall*.

DEVELOPER SIDE

Context windows fill up in **20 minutes**.

Every turn, the LLM re-sends the entire conversation. Every file read, every command output, every grep result piles up. By Friday afternoon the agent is spending more tokens carrying context than reasoning over it.

The industry response: *buy more context window*. Rate-limiting the wrong axis.

ORGANIZATION SIDE

Leaders have no **audit trail**.

A CTO sees a \$50K AI tool bill but no return. An EM sees engineers stuck without knowing. A CISO sees prompts touching secret-shaped paths without consent. Their AI tools ship faster than their evidence trail.

No incumbent observability tool sits at this layer. The category is open.

Sandbox the tool. *Index the output.* Stop sending raw bytes.

The **context-mode** plugin intercepts every tool call (bash, file reads, grep, fetch, eval) and writes the output to a local SQLite FTS5 store. The agent gets a few hundred bytes plus a search handle. When it needs the raw output, it queries the store.

```
// Without context-mode $ grep -r "useAuth" src/ → 4,200 lines streamed into the LLM context window // With context-mode $ grep -r "useAuth" src/ →  
sandboxed. 4,200 results indexed. handle grep-2026-06-12-7a8f → 327 bytes sent to the model instead of 142,000 bytes. // Agent queries the index on  
demand claude → ctx_search("useAuth") limit=10 → 10 most-relevant matches
```

98%

CONTEXT WINDOW SAVED

15

AI CODING PLATFORMS

ELv2

OPEN SOURCE

0 KB

DATA LEAVES YOUR MACHINE

Context Mode *Insight.*

• LIVE • SELF-SERVE • \$20 / SEAT / MONTH

The same plugin optionally forwards **structural events** to the org (tool names, error counts, file paths). **Never source code. Never prompts. Never model responses.** Insight runs **222 detection patterns** server-side and serves answers via MCP, role-narrowed to six personas.

CTO

What is our AI investment returning?

EM

Who is stuck on my teams this week?

IC

How am I doing vs my team?

CISO

Any high-severity AI risks?

FINOPS

Where is our AI spend going?

DEVOPS

Is commit quality healthy?

Honest framing. Insight launched in Q2 2026. *Zero paid customers, zero active org users today.* The product ships and is live. The seed converts the 290K-developer OSS base into the first paying organizations. **That conversion is the bet.**

The plugin grew to *290,000+* developers with zero paid acquisition.

←

290,000+

DEVELOPERS

17,200+

GITHUB STARS

#1

ON HACKER NEWS

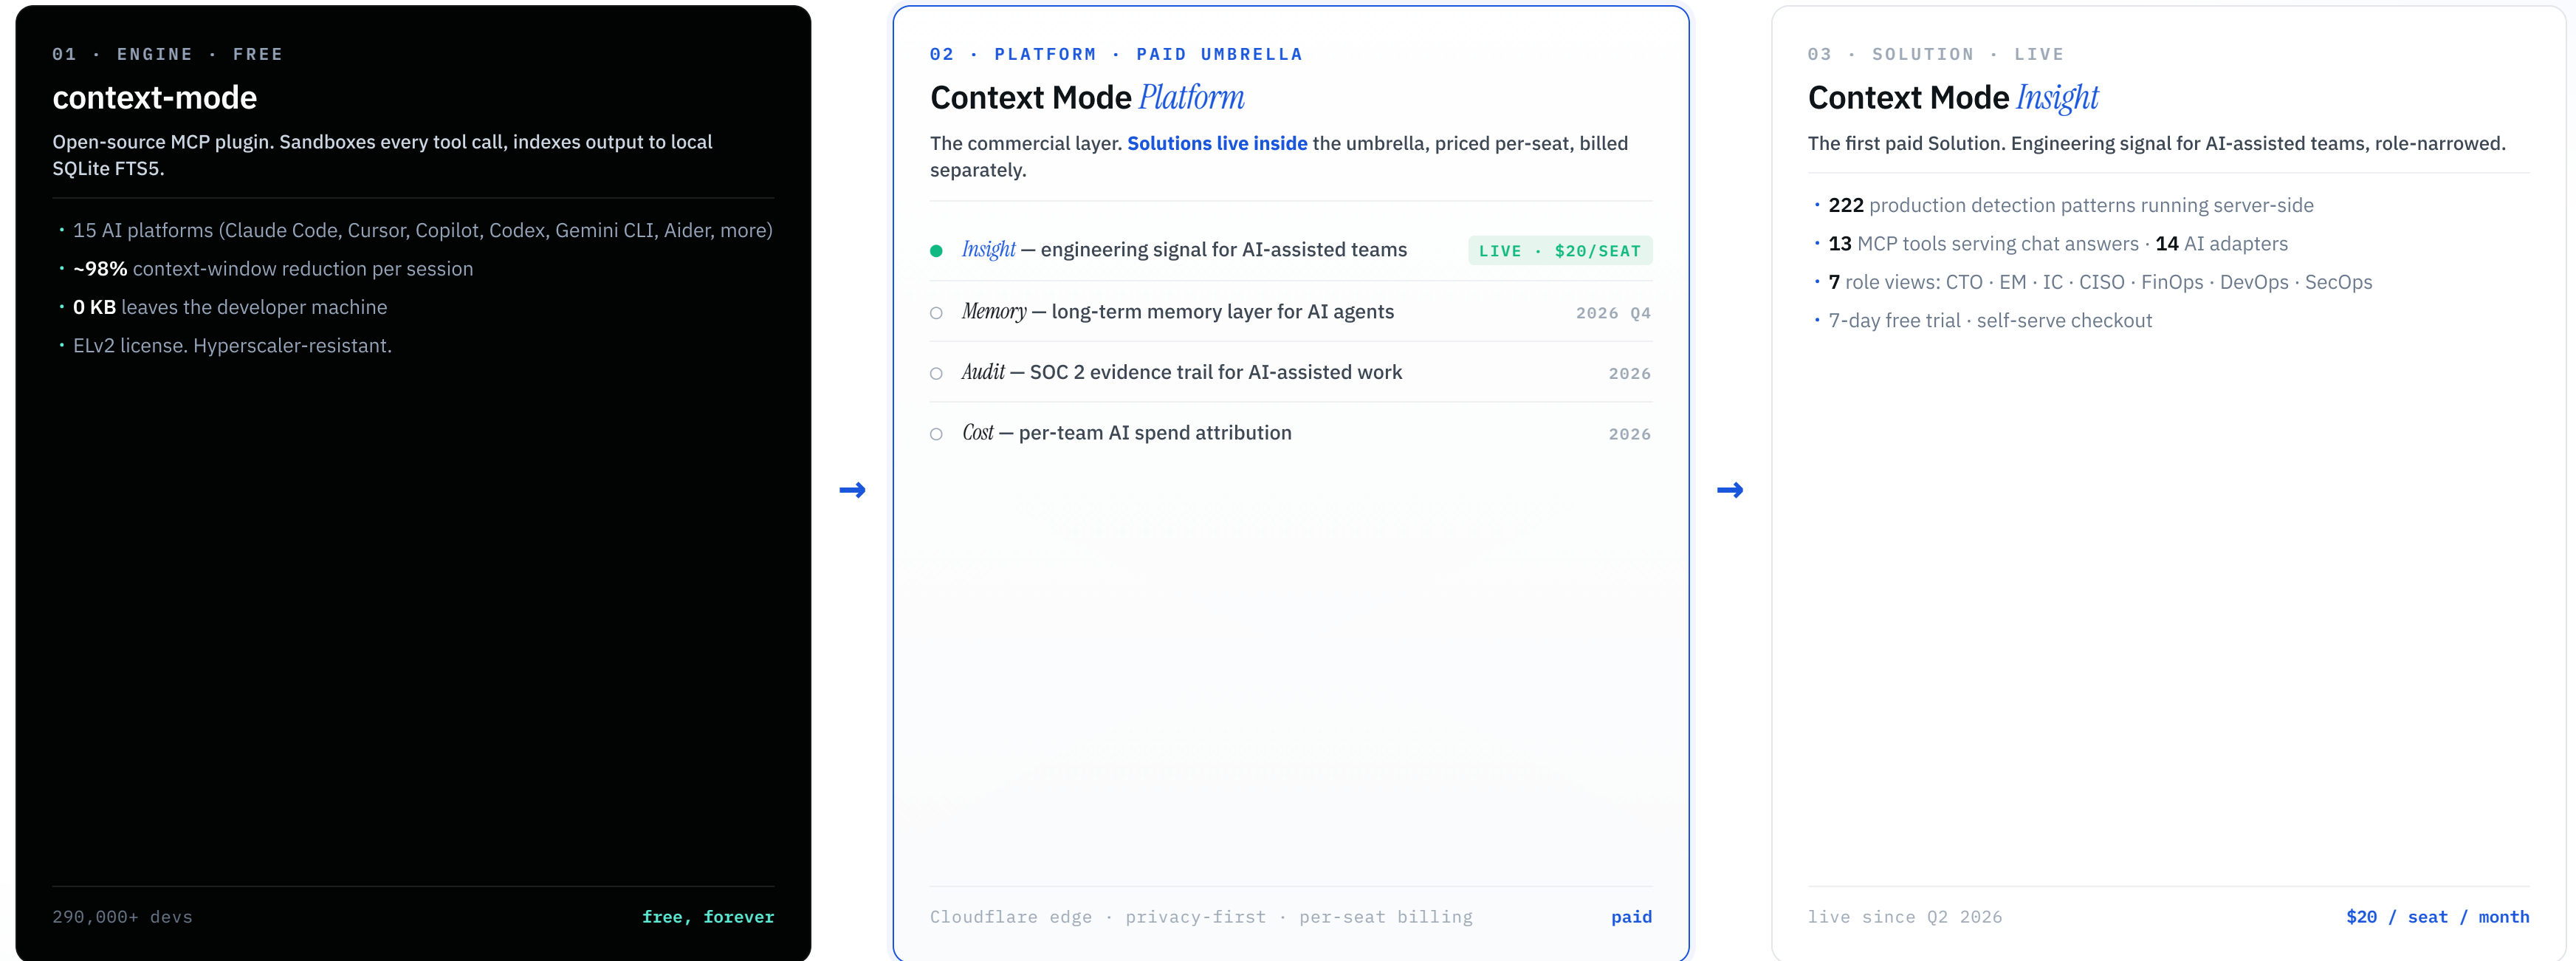
→

15

AI PLATFORMS

Organic, developer-to-developer. No outbound. No ads. The HN run, GitHub virality, and 15 tool communities did the work.

One engine. *A bundle of Solutions.*



The bundle thesis. One free engine drives adoption. The paid Platform monetizes the same installed base across multiple Solutions over time. Each Solution sells independently; together they compound into one ARR line per organization.

A *\$12.8B* tool layer with no observability incumbent.

\$1.2B

Insight TAM. AI coding assistants are a **\$12.8B** market in 2026 (projected \$30B by 2032). At ~10% observability take-rate, that is **\$1.2B**. Cursor alone reached **\$3B ARR** with 1M+ paid seats; GitHub Copilot has **4.7M paid subscribers**. Insight at \$20/seat is the lens for that installed base.

MOAT 01

Distribution already done

290K developers running the plugin. The Insight upsell is one opt-in switch away. The bridge is built; the seed funds someone to walk people across it.

MOAT 02

222 named decisions

The pattern catalog is not 222 lines of code. It is 222 product decisions about what to surface to which role. A clone needs to rebuild every one.

MOAT 03

Privacy posture

No source code, no prompts, no responses ever leave the developer machine. The only enterprise-viable posture. Incumbents trained on send-everything-to-cloud cannot retrofit it.

MOAT 04

ELv2 license

Source-available, hyperscaler-resistant. Cannot be commercialized by AWS / Azure / GCP. Same shield Elastic, MongoDB, Redis used to defend their commercial layer.

Solo founder shipped this. *Now two founding hires.*

TEAM TODAY

B. Mert *Köseoğlu*

Solo founder. Built the engine, the platform, and Insight end to end. Full-time. Based in Türkiye. New Delaware C-Corp on close.

- 290,000+ developer plugin, ELv2 OSS
- Cloudflare Workers + D1 platform, \$0 fixed infra
- 222-pattern Insight catalog · 13 MCP tools · 7 personas
- 15 AI coding platforms supported on day one

THE ASK

\$750K *SAFE*

on \$12M post-money cap

18 months minimum runway, 24 months with buffer. Two immediate founding hires: a commercial lead to convert the OSS base and a deployed engineer to walk Insight into the first design-partner orgs.

- ~\$120K founder compensation
- ~\$240K Sales+CS+Marketing+DevRel hire (18 mo)
- ~\$90K Founding Forward Deployed Engineer (18 mo · equity-weighted)
- ~\$50K growth · ad credits · sponsorships · content
- ~\$250K seed-bridge buffer · runway extension to 24 mo

WHAT WE WILL NOT TAKE

Strategic money from any AI coding agent vendor. **Board seats** at pre-seed. **Flips** from ELv2 to permissive.