

01 · PRODUCT ONE-PAGER

A language. A runtime. A product that proves it.

Agent behavior today is folklore — benchmarks, anecdotes, and vibes. Enterprises need agents that work reliably, transparently, and under control.

02 · THE PLATFORM

- Bounded work.** Every task carries a contract: scope, non-goals, validation plan, verified completion. Agents can't wander; "done" means proven done.
- Auditable cognition.** Every prompt, tool call, and state transition recorded as a durable, replayable record. Any run can be reconstructed and inspected.
- Operational trust.** The Freedom Gate checkpoints every consequential action: allow, refuse, or route to a human — with evidence recorded either way.
- Continuous Adversarial Verification.** Red-team agents attack what blue-team agents defend; nothing counts as fixed until its replay bundle fails to reproduce.

03 · THE PRODUCT: CODEFRIEND

AI-native software engineering built on ADL, and the platform's proving ground. Agent teams deliver code through the Cognitive SDLC: every change bound to an issue, validated, and reviewed from four independent perspectives — correctness, security, adversarial, and policy compliance. The runtime keeps the evidence; your engineers keep authority.

Architecture-first, not autocomplete · whole applications, not files · the first AI-centric SDLC. Status: in development — early access via partners@agent-logic.ai.

PUBLIC COMMITS

4,369

RELEASES

24

CAV RUN WP12-4914

6/6 refused

MVP TARGET

v0.95

04 · WHY NOW

Cognition is a commodity; control is not. Agents are gaining tool authority over repositories, infrastructure, and credentials — and enterprise trust requires evidence, not fluent output. Analysts forecast the AI agents market growing from \$7.8B (2025) to \$52.6B by 2030. Every one of those agents will need a control plane.