

01 · COMPANY DESCRIPTION

Agent Logic, Inc.

One line: Agent Logic builds the deterministic control plane for enterprise AI agents.

02 · THE PROBLEM

Enterprises are moving from AI that answers questions to AI that performs work. But today's agents are stateless, unauditable, and ungoverned — they live in scattered prompts and glue code that no one can inspect, version, or replay. In environments where trust, repeatability, and operational control matter, "the agent did something" is not an acceptable answer. You can't govern what you can't verify.

FOUNDED

2026

PUBLIC COMMITS

4,369

CORE RUNTIME

Rust · v0.91.8

HQ

Silicon Valley

03 · WHAT WE BUILD

- 01 ADL — the language.** Defines agent systems explicitly: agents, tools, permissions, workflows, and validation rules, declared in one specification a human can read and a runtime can enforce.
- 02 ADL Runtime — the control plane.** Executes what the language declares: enforces policy, gates tool use, and records every prompt, tool call, and state transition as a durable, replayable record. Every task carries its own contract — "done" means proven done, not asserted done.
- 03 CodeFriend — the proving ground.** Applies the platform to software engineering through the Cognitive SDLC: every change bound to an issue, validated, and reviewed from four independent perspectives. In development, with early access through the design-partner program.

04 · HOW WE'RE DIFFERENT

Architecture before autonomy. Governance is enforced by the runtime, not requested politely in a system prompt. Consequential actions pass through policy-and-evidence checkpoints that can allow, refuse, or route to a human.

Auditable cognition: plans, decisions, reviews, and outputs become evidence surfaces, not chat residue. Any run can be reconstructed and inspected long after it happened. Continuous Adversarial Verification (CAV) turns security into a loop: governed red-team agents hunt for exploits in the systems blue-team agents defend, and no vulnerability counts as fixed until its replay bundle fails to reproduce.

The work is research-grounded: formal foundations published ("A Mathematical Theory of General Intelligence," June 2026, DOI: 10.5281/zenodo.20738007), the C-SDLC preprint in review, and the Universal Tool Schema (UTS) proposed as an open standard for provider-neutral tool interfaces.

05 · LONG-TERM VISION

The next decade of enterprise AI will require a common execution substrate analogous to an operating system: infrastructure for trustworthy coordination, persistence, governance, and collaboration across many specialized agents. Agent Logic is building that foundation.