

01 · INTERNAL REFERENCE

Standard responses

Copy-paste source for applications and outreach. Keep consistent with the company description and pitch deck.

02 · COMPANY DESCRIPTIONS BY LENGTH**One-liner (~15 words)**

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

~50 words

Agent Logic builds the deterministic control plane for enterprise AI agents. Our platform, ADL, declares agent systems in an enforceable specification and records every prompt, tool call, and state transition as replayable evidence. Our first product, CodeFriend, applies this to AI-native software engineering. Founded 2026, Silicon Valley.

~100 words

Agent Logic builds the deterministic control plane for enterprise AI agents. Our platform, ADL (the Agent Design Language), declares agent systems — agents, tools, permissions, workflows, and validation rules — in one enforceable specification, and its runtime records every prompt, tool call, and state transition as replayable evidence. Agents can't wander: every task carries a contract. Our first product, CodeFriend, applies this to AI-native software engineering — agent teams delivering code through the Cognitive SDLC, reviewed from four independent perspectives, while engineers keep authority.

03 · COMMON APPLICATION ANSWERS**What problem do you solve?**

Enterprises are moving from AI that answers questions to AI that performs work — but agent behavior today is unauditable folklore. We make agent work bounded, auditable, and governed: every task carries a contract, every action leaves evidence, and human authority is enforced by the runtime rather than requested in a prompt.

What stage are you at?

Company launched publicly in July 2026. ADL platform built in the open: 4,369 public commits, 24 releases, Rust core at v0.91.8 (MVP at v0.95); formal foundations published (DOI: 10.5281/zenodo.20738007). CodeFriend in development with an open design-partner program. Pre-revenue; raising a \$3M seed while pursuing non-dilutive funding.

Why now?

Model capability has outrun operational control. Every enterprise piloting agents is discovering the same gap: they can make agents act, but cannot prove, replay, or govern what agents did. The control plane is the missing infrastructure layer.

What makes your team qualified?

Founder Daniel Austin (Ph.D, Physics, MIT) has spent three decades architecting trusted planetary-scale systems: Chief Systems Architect at Visa, Fellow at Walmart, PayPal Central Architecture Team, AWS, Sun, Yahoo!. Six startups, two acquisitions, two IPOs — and the researcher behind the platform's formal foundations.

How will you use credits? (credits programs)

Cloud credits directly fund ADL Runtime development and CodeFriend's design-partner deployments: model inference for multi-agent workflows, CI infrastructure for Continuous Adversarial Verification runs, and replay/evidence storage. Credits are non-dilutive capital that extends runway while we complete the MVP and convert design partners.

What is your long-term vision?

A common execution substrate for enterprise AI — an operating-system layer enabling trustworthy coordination, persistence, governance, and collaboration across many specialized agents. CodeFriend proves the model; ADL becomes the infrastructure.

04 · OUTREACH TEMPLATES**Investor intro (subject: Agent Logic — deterministic control plane for enterprise agents, \$3M seed)**

Hi [Name] — I'm Daniel Austin, founder of Agent Logic. We build the control plane enterprises need to run AI agents in production: a language for declaring agent systems and a runtime that enforces policy and records every action as replayable evidence. Our first product, CodeFriend, applies this to enterprise software engineering. Before this I was Chief Systems Architect at Visa and a Fellow at Walmart; I've founded six startups, with two acquisitions and two IPOs. We're raising a \$3M seed — 18 months to CodeFriend GA, five design partners, and first ARR under contract. Deck: agent-logic.ai/pitch. Would a 20-minute call in the next two weeks make sense?

Design partner (subject: Governed AI engineering — design partner invitation)

Hi [Name] — teams adopting AI coding agents keep hitting the same wall: the code arrives, but the evidence doesn't. CodeFriend, built on our ADL runtime, delivers agent-written changes bound to issues, validated, and reviewed from four independent perspectives — with every step recorded and replayable. Your engineers keep authority; the runtime keeps the evidence. We're selecting a small group of design partners ahead of launch. Partners get early access and direct influence on the roadmap. Interested in a short call?

Grant inquiry (subject: Trustworthy AI infrastructure — [Program] inquiry)

Hi [Name] — Agent Logic builds open, auditable infrastructure for AI agents: replayable execution records, enforced human authority over consequential actions, and the Universal Tool Schema, our proposed open standard for provider-neutral tool interfaces. Our goal is an AI ecosystem where agent behavior is verifiable rather than folkloric — which we believe aligns closely with [Program]'s mission. Is [Program] currently accepting proposals in this area?