01 The Foundry

Deciding what to build before an autonomous factory builds it

PERIOD
June – September 2026

STATUS
In production use inside Jiyu

MY ROLE
Designed, directed the build, Business Authority on every run

IN ONE LINE
10 ideas · 5 killed · 1 graduated · 0 advanced without pre-registered bars

Steven Fellows · Co-Founder, AI Strategy & Operations, Jiyu Technology Inc · stevenpfellows@gmail.com · 206-483-4426 · stevenpfellows.com · One-page extract from the AI Transformation Portfolio; figures are from company records.

BUSINESS PROBLEM

A two-founder company with more ideas than capacity and no neutral way to choose among them. Every idea arrived with an enthusiastic sponsor, usually one of us. AI research could be made to say almost anything. And the cost of a wrong "yes" was months of build time in an autonomous factory that would happily build it. We needed a decision process that would survive our own optimism.

SYSTEM AND APPROACH

The Foundry is a seven-stage pipeline (Intake, Framing, Research, Scoring, Verdict, Conviction, Graduation) with exactly two human gates. At Gate 1 the Business Authority signs the decision bars an idea must clear before any research runs: a score floor (default 65 of 100, often raised to 70), the critical dimensions where low confidence is disqualifying, and hard gates such as legality, capital cap, and operating load. Research fans out to five to eight agents across a twelve-dimension rubric; a mandatory Bear Memo argues the kill case; a weighted scorecard produces a score and a recommendation. At Gate 2 the Business Authority reads the dossier, the Bear Memo, and the scorecard and renders PURSUE, ITERATE, PARK, or PASS. Every verdict and every bar change is written to a Decision Log and an Idea Register, and a dashboard reachable from a phone shows what is waiting on a human.

MY ROLE

I directed the specification (drafted by an AI assistant at my direction), set the governance rules, and served as Business Authority on every run: signed the bars at Gate 1, read the research, rendered the verdicts at Gate 2. Two design decisions were mine specifically. When the first design asked me to claim an "upside tier" for an idea at intake, I refused ("that's what the research is for"), so tiers are now assigned from evidence at scoring. When the system offered a static default bar to sign, I called it rubber-stamping; the redesign shipped 2026-06-16 tunes bars per idea and shows the change as a diff for sign-off. The core library, Cowork skill, and dashboard were built by AI coding agents under the company's gated delivery process; my co-founder operates the hosts. I did not write code.

OUTCOME

Ten distinct ideas (eleven runs) between June and September 2026. Five closed PASS, one graduated PURSUE (Ask Jiyu, 81 of 100, all gates passed, built as a factory test article at zero cash cost), three scored PASS and await verdict, and the rest are in research. Scores ranged from 33 to 81. No idea advanced without meeting bars signed before its research ran. A Tesla rideshare-rental idea scored 46 against a signed 70 floor after research found the leased variant contractually void and the owned variant negative before insurance. A copy-trading proposal scored 34 against a 70 floor; I ruled it a signup scheme rather than a business, and the evidence agreed. On that last run the system caught its own failure: a dashboard button had written a stub framing brief with default bars, so my first sign-off was on the wrong document. The register was rewound the same day and I re-signed the correct bars before research ran.

ENTERPRISE RELEVANCE

This is the AI use-case intake and prioritization problem every enterprise program has: a flood of proposals, sponsors with incentives, research that can be produced on demand, and a portfolio committee that cannot tell rigor from enthusiasm. The Foundry's shape transfers directly: pre-registered success criteria per use case, AI-assisted research with a mandatory red-team memo, a weighted scorecard the committee can inspect, a named human who signs the criteria before the study and renders the verdict after, and a decision log that survives turnover and audit.