Ethan Yi Sang

AI Product Manager & Builder · Product, design and code, from idea to App Store

Open to any location and time zone — remote or relocation · rco13.madlax@gmail.com · +86 132 5822 9091 · ethanflow.com · github.com/Ethan-YS · linkedin.com/in/ethan-ys

SUMMARY

I build AI products end to end, hands-on: I find the gap, design the product, build it myself with Claude Code and Cursor, ship it to real users and keep iterating from how it is used. Over the past year and more I have independently built nearly 30 AI products and tools: 2 paid products on sale (real Stripe revenue), 4 apps on the App Store (Mac + iPhone), 3 live platforms, 6 open-source projects (project-brain, 170+ GitHub stars), and more than ten agent systems I run every day. From spotting the gap to launch: one day at the fastest, three and a half months at the slowest. I make the product calls too: scope, boundaries, what to cut, and when to redefine a product. Grounded in 10+ years of product and interface design for B2B, government and consumer products.

Stack: Swift / SwiftUI (macOS, iOS), TypeScript, Python, React, Next.js, Node.js, FastAPI, Electron · SQLite, Supabase, Cloudflare Workers, REST, WebSocket · MCP (local and remote), OAuth 2.1 · Stripe, Ed25519 license activation, notarized distribution, App Store releases · on-device ML on the Apple GPU (ASR, diarization, TTS) · RAG / hybrid retrieval · built with Claude Code, Cursor, Codex CLI.

SELECTED WORK (SPROUT LABS, MY PERSONAL PRODUCT LINE)

VoxAI — AI meeting copilot product line On sale + 2 apps in stores · voxai.ethanflow.com · 2026.03–present

- Designed, built and shipped three products: the macOS meeting copilot (kickoff to public sale in 3.5 months), the Mac voice-input edition (kickoff to the Mac App Store in 8 days; 1,946 lines of pure Swift, zero dependencies) and the iPhone recording & transcription edition on the App Store.
- Real-time meeting transcription with speaker labels, fed to any AI over local MCP; the AI pushes fact-check, risk-alert and negotiation-strategy cards into the meeting and can talk back by voice. Transcription, speaker separation and TTS run locally on the Apple GPU; data never leaves the device.
- Built the full commercialization chain in 3 days: notarized direct distribution, in-app auto-update, Ed25519 offline activation, Stripe → Cloudflare Worker; payment to license code in seconds with zero manual work.
- Product decisions: redefined the product mid-way from a transcription tool into "a conversation copilot that puts AI in the room"; call mode records two channels (microphone = me, system audio = the other side), turning speaker separation from an ML problem into a physical fact; replaced the 14-day trial with 10 free meetings before a one-time purchase, because meetings are infrequent and a day-based clock is unfair to light users.

AndCo — a social platform for people and their own AIs Live · andco.ethanflow.com + App Store · 2026.09–present

- Designed and built independently: people bring their own AI into groups, and AndCo gives those AIs identity and relationships. Web live one day after final design (8 people, 11 AIs in the first 24 hours); native SwiftUI iPhone app on the App Store five weeks later, then 5 App Store releases in 5 weeks (1.0 → 1.0.4).
- Built my own OAuth 2.1 server (PKCE, device code) + remote MCP (20 tools), tested with Claude and ChatGPT clients; added AndCo Pass (a signed, site-bound credential proving whose AI it is) and OpenID Connect sign-in.
- Iterated from real failures after launch: 4 versions of the agent wake-up client (one AI was unreachable for 19 days behind silent retries). 171 automated tests; 1,450 req/s, p99 < 320 ms at launch.

Soul Engine Next — modular personal agent In daily use · native macOS · 2025.10–present

- Native macOS agent harness (Swift / SwiftUI): memory and skills live in a user-owned, versioned folder, while Claude Code, Codex or any ACP-compatible agent plugs in through its own engine adapter, swappable even mid-conversation; multi-window with sub-agents.
- Capabilities are MCP modules whose permission manifests are enforced by a macOS sandbox; credentials stay in the Keychain, never seen by modules or the agent; human-in-the-loop approvals; overnight memory consolidation; agent-written skills ship only after hidden evals and human sign-off. 235 commits and 255 tests in 13 days; first generation built since 2025.10, now on public sale.

Taproot — authorization & audit layer for a personal AI agent Runs every day · 2026.07–present

- Built the layer my personal agent works under every day (account checks, calendar sync, job-search follow-ups): every grant needs my approval: one-time grants are single-use, pinned to one action on one target and approvable from my phone; standing grants can only be signed with Touch ID on my Mac; the agent uses API keys without ever seeing them; every grant, use and refusal is recorded, replayable and tamper-evident. 400+ automated tests hold these boundaries.
- While I'm at work it runs its own shift and splits work among sub-agents; payments and signatures always come back to me. Built to answer the first question agents face in real life: what can AI do on a person's behalf, what can't it, and what guarantees that.

Product matrix, nearly 30 in all: on sale / in stores — VoxAI Meeting Copilot, Soul Engine (first generation: local-first AI companion, soul.ethanflow.com), VoxAI Voice Input, VoxAI Mobile, Twoday (a calendar shared by a person and their own AI), AndCo iPhone app · live platforms — AndCo (web), AgentMart, Helm · open source — project-brain (170+ stars), ccvault, PromptGuard for Agents, AI Agent Trust Boundary, QuoteMind · 12 agent systems running daily (personal voice knowledge engine, people & commitments knowledge graph, hybrid-retrieval memory engine, on-device visual security agent, Tactree, JobFlow-AI and more) · 4 prototypes.

EXPERIENCE

Deli Group — AI Research Center — Senior AI Product Manager 2026.10–present · Hangzhou

- Define a meeting-scenario AI product line from zero: hands-on testing of competitor and in-house recording devices, product architecture and a staged validation plan; build the interactive prototypes myself with AI tools, so ideas are tested as working software.

Sichuan Everthese Technology — AI Product Manager (B2B, 0→1 lead) 2026.04–2026.06

- Took "Hongzhi," an AI work-efficiency SaaS for state-owned enterprises, from zero to a high-fidelity interactive React product: positioning, architecture and AI interaction design; redefined it from a "Party-building system" into AI-assisted work-efficiency software after user research at 2 state-owned enterprises.
- Designed the "goal custody" agent pattern (goal → AI proposes a visible plan → human reviews and approves → brakes at critical steps → receipt) and a four-layer permission model for one person holding several roles.

Sichuan Zhizai Shifang Technology — UI Lead / Design-System Owner (led 2) 2025.04–2025.11

- Built runnable AI-assisted prototypes that cut requirement validation from weeks to 5 days; compared a low-code route with a full AI pipeline and rolled the AI workflow out in-house (~30% more efficient R&D process).
- Built the company's Figma component & variable library from scratch and 2 Figma plugins (one bundling 7 tools); co-built program templates with frontend (design efficiency +50%). 12 government & B2B projects, including a ¥10M-scale big-data remote support center.

Independent Studio (own venture) — Product Lead 2023.09–2025.02

- Co-founded a small studio with several developers serving SMEs and government; owned each project hands-on from requirements and product solution through interaction and visual design to delivery, across wealth management, education, fresh-food logistics and fitness. Delivered the WE run fitness & health app for a Japanese client from 0 to 1.

Guizhou Yotecloud Technology — UI Design Lead (led 5) 2020.09–2023.07

- Led a 5-person design team delivering the Guizhou Nanming new-infrastructure series (local lifestyle, vocational education, business district) and the "Magic Cube" low-code platform, both launched.

Juyuan Yinhai / Chengdu Xingzheng E-Government — UI/UX Designer 2018.05–2020.04

- Sichuan Health QR Code (30M+ users, CCTV national news coverage): core team member, led the Elder-Friendly Mode. Also the Tianfu Easy-Service app and government big-data visualization.

SKILLS

Building with AI: Claude Code (skills, hooks, plugins, sub-agents), Codex CLI, Cursor, Kiro, Antigravity, Gemini CLI, Figma Make, used daily from first prototype to App Store release.

Engineering & agent architecture: agent harness (built my own multi-engine harness), MCP (several self-built servers, sandboxed modules, OAuth 2.1 remote access), multi-agent collaboration, agent evals and gating, AI memory systems, context management and long-task handoff, local-first AI, RAG and hybrid retrieval (dense + sparse RRF), recall benchmarking (Recall@5 62.5% vs 54.2%), voice AI pipeline (ASR, diarization, TTS), knowledge graphs, on-device vision.

Design: 10+ years of product, interaction and interface design; design systems (Figma components and variables, 2 self-built Figma plugins); real-time AI interfaces; prototyping in working code (React, SwiftUI); Elder-Friendly Mode for a 30M-user public app.

Product: 0→1 product definition, finding gaps and setting scope, validation with runnable prototypes, commercialization (Stripe, licensing, notarized distribution, App Store releases), user-journey funnels and growth experiments.

Agent safety: human-in-the-loop approvals and brakes at critical steps, single-use scoped grants, secret isolation, audit trails and replay, prompt-injection and overreach defense, unattended operation with human takeover.

Languages: Chinese (native), English (written: professional; spoken: conversational).

EDUCATION

B.Eng. Microelectronics, Changchun University of Science and Technology, 2010–2014 · Minors: Visual Communication Design; Principles of Human–Computer Interaction