




STARTUP GROWTH ACCELERATOR

Burn your free credits. **Ship like a 50-person team.**

Open-source agent orchestration deployed on the Azure and AI credits you already have. 10× your team's output without hiring or burning cash runway.

Free · Open source · One-click deploy Powered by Claude · Runs on Azure · Microsoft Partner



Three pains every AI startup feels

If you're an AI startup founder, you probably recognise yourself in at least one of these. Often all three.

Credits you can't burn

You won Azure credits, OpenAI tokens, Anthropic credits, partner perks. They sit unused while your runway burns. Nobody on the team has bandwidth to actually deploy them as horsepower.

Founders in everything

You're writing code, reviewing PRs, handling IT, talking to investors. Every feature waits on someone — usually you. The team is small and scaling yourself is not a strategy.

Ideas faster than the team

Users tell you what to build. Investors ask "what's next." Your backlog grows by the day. Each release feels good but the gap between what you *could* ship and what you *do* ship widens every week.

These aren't solved by hiring. You don't have the cash, the time, or the months for onboarding. They aren't solved by buying more AI tools either — you already have credits sitting unused. **The lever is to put those credits to work as an agent-orchestration pipeline that ships for you, while you stay in control.**

That pipeline is Ivy Tendril. It's free, open source, and one-click-deployable onto the Azure subscription where your credits already live. You bring the credits and the ideas. Tendril converts them into shipping code.

A growth-acceleration push you can do basically for free

As an AI startup you sit on a stack of assets you've never deployed: **Azure credits, AI tokens, partner discounts, accelerator perks.** Most expire unused. Combine them with Tendril and you get a continuously-running pipeline that converts ideas, bugs, and feature requests into production-grade code at the speed your founders *ideate* — not the speed your team types.

The simple math

MICROSOFT FOR STARTUPS FOUNDERS HUB	up to \$150,000 Azure credits
+ ANTHROPIC / OPENAI STARTUP CREDITS	\$5K-\$25K AI tokens
+ IVY TENDRIL (FREE, OPEN SOURCE)	\$0
= A MAX-CAPACITY AGENT PIPELINE RUNNING FOR	months, not minutes

Every token your agents burn is a token *you didn't pay for in cash.* Every parallel agent the orchestrator runs is throughput you didn't have to hire. **Your unused credits become the engine of a 10x productivity push that costs you nothing on the P&L.**

You already paid for the fuel. Tendril is the engine. The next four pages cover: where the credits are, where your team sits on the maturity curve, how the one-click deployment works, and the three packages we offer to get you there fast.

Where the credits live

If you haven't applied yet, start here. The programs below are free, application-based, and stack with each other. Most can be approved in days, not weeks.

RECOMMENDED · START HERE

Microsoft for Startups Founders Hub

up to \$150,000 Azure

Free Azure credits, Azure AI Foundry (incl. Anthropic models on Azure), Azure OpenAI Service, GitHub Enterprise, Visual Studio Enterprise, plus 1:1 advisor support and partner perks.

No funding required to apply. Tiered: \$1K → \$5K → \$25K → \$150K as you progress.

Apply: aka.ms/startups

AI TOKENS

Anthropic for Startups

\$5K–\$25K Claude credits

Direct API credits for Claude. Strongest path if Tendril is running Claude as the primary model. Typically requires accelerator or VC backing.

Apply: anthropic.com/startups

CLOUD + AI

AWS Activate

up to \$100K AWS

AWS credits including Bedrock (Claude, Llama, Mistral). Higher tiers gated on accelerator or VC partner introduction.

Apply: aws.amazon.com/activate

CLOUD + AI

Google for Startups Cloud

up to \$200K GCP

Google Cloud credits over 2 years, plus Gemini access. Good fit if your stack is already on GCP or if your model mix includes Gemini.

Apply: cloud.google.com/startup

Already applied? Skip ahead to page 6 — we'll deploy Tendril on the subscription where your credits live and burn them productively from day one. **Haven't applied yet?** The free workshop in Package One walks you through the Microsoft Founders Hub application together.

Where AI startups are today, and where the unicorns are heading

Steve Yegge's *8 Levels of AI-Assisted Development* (Augment Code, 2025) is the cleanest map of how teams actually use AI today. The same diagnostic applies to AI startups — with one twist: the funded competitor next door is sprinting toward level 8 while most teams are still hand-typing prompts.

LEVEL	WHAT DEVELOPERS DO
1 – 2	Autocomplete, copy-paste from a chat panel.
3	Highlight code, ask the IDE to edit it inline.
4 – 5	Let agents run; work in the terminal while agents handle async tasks.
6 – 7	Run multiple agents in parallel, but coordinate them by hand.
8	Custom orchestrator. A shared task queue, a coordinator process, and checkpointing manage many agents working together reliably. The team directs <i>infrastructure</i> , not individual agents.

Most startups plateau between levels 2 and 3. Buying a more advanced agent does not move you up. The shift to level 8 is *structural*: it requires engineering the workflow around the agents. That is exactly what Tendril is built to do.

Why this matters for an AI startup specifically

At levels 2–3 you get *individual* productivity — one founder + one agent ships a bit faster. At level 8 the whole team's throughput **compounds**: the orchestrator coordinates agents across people, repos, and tickets, and it accumulates institutional knowledge about how *your* codebase wants to be worked on. **That is the 2x–10x output gap that funded competitors will use to outrun you** — not better models, but better orchestration.

Ivy Tendril is built to take an AI startup from level 2–3 to level 8 on the credits you already have — with a free workshop to wire it up.

A max-capacity Tendril server, deployed in minutes

One click puts a high-capacity Tendril orchestrator on your Azure subscription — GPU-backed worker pool, persistent memory, autoscaling, and a coordinator that runs **10+ agents in parallel**. Powered by Claude by default. Funded by the credits you already have.

DEPLOY

Deploy in minutes

A single command spins up the full Tendril stack on your Azure subscription. No DevOps tickets, no IT dependency, no months of infra setup.

CAPACITY

Max-capacity sizing

Provisioned for parallel agents from day one. The coordinator schedules across a worker pool that scales up when work piles up and down when it doesn't.

MODEL

Claude-powered, model-agnostic

Ships with Claude as the default model. Swap to Codex, Gemini, or a local model with one config change — route specific tasks to specific models.

MEMORY

Persistent institutional memory

The orchestrator learns *your* codebase over time. Every run improves the next. Patterns, conventions, edge cases — remembered, not relearned.

FREE. OPEN SOURCE. ONE-CLICK
DEPLOY.

PLUS: A **FREE** AI WORKSHOP TO WIRE IT INTO
YOUR STACK.

You can click it yourself. The source is at github.com/Ivy-Interactive/Ivy-Tendrill and the docs are at tendrill.ivy.app. If you'd rather not click alone, every package on the next three pages includes a **free 90-minute workshop** with an Ivy engineer to deploy it onto your subscription, wire it to your repo, and put your first production agent live.

PACKAGE ONE · ONBOARD

Free AI Workshop

Get your engine running — on us.

A free 90-minute working session where an Ivy engineer deploys Tendril onto your Azure subscription, wires it to your repo, and puts your first production agent live alongside you.

YOU LEAVE THE SESSION WITH

- **Tendril deployed** on your Azure subscription, sized for max capacity, burning the credits you already have.
- **Your first custom agent** wired to your repo — code review, test generation, or whichever task is the highest-leverage offload for your stack.
- **A shortlist of next agents** to add — the 3–5 highest-leverage tasks specific to *your* codebase, prioritised by effort vs throughput gain.
- **Help applying for credits**, if you haven't yet — we walk through the Microsoft Founders Hub application together.
- **Direct access to our Discord** for follow-up questions and pattern sharing with other startups in the program.

Free · Open source · One-click deploy. No fees. No contract. No "trial period." We do this because we want AI startups running Tendril, and the fastest way to get there is to put an engineer in the call with you.

BEST FOR

Any AI startup that wants Tendril running on their subscription without burning a weekend on it. No prerequisites.

FORMAT

90 minutes, video call. We screen-share the deployment with you, live.

PRICING

Free.
Book at
renco@ivy.app.

PACKAGE TWO · ACCELERATE · MOST POPULAR

Growth Accelerator

Idea → production code in 2–3 weeks.

We embed with your team for one focused sprint and build the full ideas-to-shipping pipeline. Your founders specify; agents code; humans gate. By the end you're operating at level 6+ with credits as the fuel and Tendril as the engine.

WE DELIVER, ALONGSIDE YOUR TEAM

- ▶ **Custom orchestration setup** tuned to your repo, your stack, your conventions. Planning agent → coding agent → review/test/security agents → human-gated PR.
- ▶ **3–5 production-grade agents** shipped live: code review, test generation, security audit, doc updates, ticket grooming — whichever set returns the most for your team.
- ▶ **Founder-facing ideation interface** so anyone on the team can drop in a feature request or bug and get scoped, agent-ready work back in minutes.
- ▶ **Persistent memory configured** so the orchestrator learns your codebase and gets sharper every week instead of starting from zero.
- ▶ **Free AI workshop included** — this is the first thing we do.

Free · Open source · One-click deploy. The software stays yours, MIT-style. We build with you, not on top of you.

BEST FOR

Funded AI startups (Pre-seed to Series A) ready to convert dormant credits into shipping output. Teams of 2–15 engineers.

FORMAT

2–3 week embedded sprint. Daily standups with our engineer; weekly steering with founders.

PRICING

We share the upside. Credit-sharing, milestone-based, or success-fee structures.
renco@ivy.app

PACKAGE THREE · COMPOUND

Full Co-Build

We embed; you scale 10x.

For AI startups going all-in on agentic engineering as a competitive moat. Ivy embeds with your team for the long haul — orchestration, custom agents, training, governance, ongoing improvement. We build alongside you until your team is at level 8 and operating without us.

WE DELIVER, OVER THE FULL ENGAGEMENT

- ▶ **Long-term embedded partnership** — one or more Ivy engineers working alongside your team, not delivering and disappearing.
- ▶ **Custom agent development** for your specific domain — the agents your competitors don't have because they're built around *your* codebase, customers, and pipeline.
- ▶ **Training and playbooks** so your team operates the orchestrator confidently and keeps moving up the curve after we leave.
- ▶ **Quarterly reviews**, continuous improvement, model migration as the frontier evolves, governance for compliance and audit.
- ▶ **Free AI workshop included** — same as every package.

Free · Open source · One-click deploy. The orchestrator is yours forever. The relationship is the engagement.

BEST FOR

Series A+ AI startups treating agent orchestration as a strategic moat, not a productivity tweak.

FORMAT

Quarterly engagement, renewable. Embedded engineers, scoped to your team size and ambition.

PRICING

Custom-scoped. Credit-sharing, success fees, equity, or hybrid.
renco@ivy.app

TALK TO US

Let's get started.

Whether you've already won your credits or you haven't applied yet, we'll meet you where you are. Book a free 90-minute workshop and we'll deploy Tendril on your subscription together.

EMAIL

renco@ivy.app

WEBSITE

ivy.app

DOCUMENTATION

tendril.ivy.app

SOURCE CODE

github.com/Ivy-Interactive/Ivy-Tendril

COMMUNITY

[Discord](#)

OPEN SOURCE

Free to use

WHY WE DO THIS

If you grow, we grow.

Ivy is a Microsoft Partner. We want to be early in the next unicorns. Bring your credits, your ideas, and your ambition — we'll bring the orchestration, the engineers, and a flexible partnership structure that fits your stage. Help us help you ship.

