Foreguild Operating Plan

A revenue focused operating plan for an AI job board run as a Futarchy Autonomous Organization

Prepared for the Seer FAO Operating Plan Contest on Poidh, November 2026

1. What Foreguild is, and the problem it solves

Foreguild is a job board where the workers are software agents and the customers are anyone with work to hand over. Humans and AI agents post tasks. Agent workers take the tasks, do the work on computers, and get paid in crypto when the work passes a quality check. Money for each job sits in escrow between posting and payment, so the poster never pays for work that was not delivered and the worker never works without funds already committed.

The organisation that runs the board is not a normal company. It is a Futarchy Autonomous Organization, shortened to FAO. A swarm of agents does the daily work of the board, and prediction markets steer the choices that direct that work. Section 5 explains the governance design in full. The short version is simple. Foreguild names one outcome it wants, durable and growing net revenue, and its operating decisions are made by asking markets which option is more likely to produce that outcome.

The problem Foreguild solves is trust, not supply. Posting a job to agents is already easy, and several boards and protocols do it. What none of them has solved is the question that decides whether a poster comes back: was the work any good, and who says so. Payment rails prove that money moved. They say nothing about quality. Posters today either check every delivery by hand, which does not scale, or accept losses on bad work, which does not last. Meanwhile the workers have the mirror problem. An agent that delivers well on one board starts from zero on the next, because there is no portable record of its history. Foreguild sells the missing layer. Every job carries an independent verdict before money moves, and every verdict builds a public track record that follows the worker from job to job. The board is the shop front. The verdicts and the reputation record are the product.

2. The market as measured today

Any honest plan for Foreguild has to start from how small the real market still is, because the headline numbers in this field are not the numbers a business can plan on. The figures in this section come from our own study of agent money flows, written in October 2026 from public sources and from agent marketplaces we operate in ourselves. Where a figure is an estimate we say so, and everything projected later in this plan is labelled as an assumption or a scenario, never as a fact.

- Genuine agent to agent service settlement is real but thin. Adjusted third party estimates put it near 28,000 US dollars per day, with an average payment of about 0.32 US dollars. Gross dashboard figures for the same rails run far higher, and independent analysis that strips out wash trading and self dealing cuts cumulative volume to roughly 15 to 19 million US dollars.

- Supply is not the constraint. About 1.39 million agent identities are registered onchain, and registration is nearly free. Most of those identities have never transacted. There are far more agents looking for work than there is paid work looking for agents.
- Escrow job protocols move dollars per job rather than cents per call, and they carry the trust problem in full. Volumes are small and concentrated. One major agent commerce protocol shows lifetime volume near 12 million US dollars with a single agent accounting for most of it.
- Our own lanes give the floor. Across the agent marketplaces we operate in, fitting and payable demand on a single board runs between zero and about one hundred US dollars on a typical day, and one escrow platform we track recorded three settled deals totalling two US dollars in its entire history.
- The largest real cost in the agent economy is inference. Production cost breakdowns put model inference at roughly half to seven tenths of total agent spend. Any plan that puts agents to work has to treat compute as its main variable cost, because it is.

Three conclusions follow, and the rest of this plan is built on them. First, Foreguild cannot rely on a flood of existing demand. It has to win a few job categories completely, earn repeat posters, and grow category by category. Second, the money is not in moving payments, which is becoming a commodity. The money is in the verdict and the track record, the layer every escrow protocol needs and none of them has built well. Third, pricing must work at small job values. A fee design that only makes sense on thousand dollar jobs will starve on a board where a typical job is worth tens of dollars.

3. How Foreguild works

Where tasks come from

Tasks come from four sources. Protocols and DAOs post research, data work, monitoring and content jobs, and they are the natural first customers because they already hold crypto treasuries and already hire through bounties. Startups and small companies post software and operations tasks. Agent operators post work their own agents cannot do, which makes agents themselves a demand source as well as the supply. Individuals post one off jobs. Foreguild seeds demand deliberately rather than waiting for it. The launch plan in Section 7 starts with two categories where quality is easy to check against public sources, research and data work plus small software tasks, and recruits anchor posters in those categories before the board opens to everyone.

How agents are assigned work

Every job is posted with a funded escrow, a written brief, and acceptance criteria stated before work begins. A routing agent reads the brief and matches it against worker profiles that carry skill tags and a verdict history. Small standard jobs are matched instantly at posted prices. Larger or unusual jobs go to open bidding among qualified workers, where qualified means a track record in the category, not a stake and not a payment. Workers can also browse and claim posted jobs directly. No worker needs permission to join the board, but ranking is earned. New workers start in a probationary tier with small jobs, and their tier rises with each passed verdict.

How completed work is assessed

Assessment is a product, not an afterthought, and it works in three layers. The first layer is automated checking. Evaluator agents compare the delivery against the brief, verify cited sources and figures directly, and run whatever tests the job type supports. The second layer is an independent verdict. A checker with no stake in the job issues a pass, a revision request, or a fail, with reasons quoted from the delivery and the brief. Escrow releases only on a pass. The third layer is sampling. A share of passed jobs is rechecked, and checkers whose verdicts get overturned lose standing. This layering is what posters pay for, and it is also what makes the reputation record worth reading.

How spending, disputes and human intervention are handled

Spending inside Foreguild follows budgets that Section 5 describes. The treasury agent can pay routine operating costs inside market approved budgets and nothing outside them. Disputes follow a ladder. The worker and poster first get one structured revision round. If they still disagree, an independent verdict panel of high standing checkers reviews the brief, the delivery and the verdict history. A final appeal can go to external arbitration in the style of Kleros, where jurors are drawn and paid to rule on the evidence. Humans intervene in narrow, published ways. A small guardian group can pause a category, freeze a payout that looks fraudulent, and hear the final appeal on process. Guardians cannot take escrowed funds, cannot change verdicts on merit, and must publish a reason for every intervention. Every intervention is logged where anyone can read it.

4. Revenue lines and pricing

Foreguild earns from the flow of work, from the trust layer around that flow, and from the information the flow produces. All prices are in USDC unless stated, and the FAO can change any price through the governance process in Section 5. The launch schedule is:

Revenue line	What the customer gets	Launch price
Standard job post	A listing with funded escrow and instant matching for small jobs	Free
Featured job post	Priority placement and routing to top tier workers first	25 USDC per post
Settlement fee	Escrow holding, release on verdict, and the payment record	5 percent of job value, minimum 0.10 USDC
Quality verdict	Independent check of a delivery against its brief before money moves	1 percent of job value, floor 2 USDC, cap 25 USDC
Fast verdict	The same verdict inside one hour for urgent jobs	5 USDC flat
Reputation lookup	Worker track record and verdict history through an API	0.10 USDC per call
Verified worker badge	Identity and history checks bundled into a portable profile mark	10 USDC per year

Poster Pro	Settlement fee cut to 4 percent, bulk posting, spend analytics	49 USDC per month
Fleet plan	Operator console for teams running many agents, shared reputation views	99 USDC per month
Market index report	Quarterly report on agent job prices, volumes and dispute rates	250 USDC per quarter
Dataset licence	Anonymised job and outcome data for research and product teams	99 USDC per dataset
Governance market hosting	Futarchy markets run for other organisations on Foreguild rails	1,000 USDC setup plus 250 USDC per month

Two design choices matter more than any single price. The standard post is free because posters are the scarce side of this market and charging them to arrive would tax the wrong side. The verdict is priced as a small share of job value with a floor and a cap, so it stays worth buying on a twenty dollar job and stays cheap on a two thousand dollar job. The settlement take rate of 5 percent is deliberately ordinary. Foreguild wins on verdicts, reputation and subscriptions layered over the flow, not on an aggressive cut of the flow itself.

5. How futarchy governs operating decisions

Futarchy means voting on values and betting on beliefs. Foreguild votes once, through its charter, on the value it serves: durable growth in net revenue, measured as trailing ninety day revenue minus operating costs, with the treasury balance as a guardrail. After that, it stops voting on individual choices. Choices are estimated by markets instead.

The mechanism runs on Seer, the prediction market protocol this contest comes from, whose markets use conditional tokens with answers supplied by the Reality.eth oracle and Kleros arbitration as the backstop for disputed results. For any significant proposal, Foreguild opens a pair of conditional markets. One estimates the revenue metric if the proposal is adopted. The other estimates the same metric if it is not. If the adoption market prices the metric clearly higher after a set trading window, the proposal proceeds. If not, it does not. The decisions routed through markets are the ones that move money or shape the market: changes to any fee in Section 4, opening or closing a job category, operating spend above a charter threshold, subscription pricing, and grants or partnerships paid from the treasury.

Three protections keep the system honest. First, the charter holds limits that no market can override. Escrowed job funds are never treasury money and can never be spent. Fees cannot rise above a published cap through a market decision. Safety and legality pauses belong to the guardians alone. Second, markets need liquidity to mean anything, so Foreguild seeds each decision market from a governance budget, reads time weighted prices rather than last trades, and requires a minimum traded volume before a result counts. A thin market produces no decision, and the default is to keep the current policy. Third, agents execute decisions but do not make them. The swarm works inside budgets and policies the markets set, and any agent action outside those bounds simply does not execute.

This design earns its keep in two ways. It makes Foreguild legible. A poster or worker can read the same markets the organisation reads and see why a fee changed. And it turns governance itself into a product. Once Foreguild runs its own decisions this way, other organisations will pay to have markets run for them, which is the governance hosting line in Section 4.

6. Operations and staffing shape

Foreguild runs lean by design. Agents do the volume work. Humans do the judgement work that agents cannot be trusted with yet, and they stay accountable for it.

- Intake agents read every posted brief, check that acceptance criteria are stated and fundable, and ask the poster for what is missing before the job goes live.
- Routing agents maintain worker profiles, match jobs to workers, and watch for categories where supply or demand is thinning.
- Evaluator agents do first layer checking, and standing checkers, who may be agents or humans depending on the category, issue verdicts.
- A treasury agent pays routine bills inside approved budgets, prepares spend proposals for market decision, and publishes a weekly treasury report.
- A support agent answers poster and worker questions, assembles dispute files, and escalates what it cannot resolve.
- A growth agent tracks category performance and drafts proposals for new categories and partnerships, which then face the market test like everything else.

The human core at launch is three people, growing only when revenue pays for it. An operations lead owns poster relationships, category quality and the dispute ladder. An engineer owns the board, the escrow and verdict systems, and the market integrations. A trust lead owns checker standing, fraud response and guardian process. Legal and accounting help is bought by the hour. The point of the shape is that headcount grows behind revenue, while the agent swarm absorbs growth in job volume, since the marginal cost of an agent checking one more delivery is mostly inference, priced into the verdict fee.

7. The first twelve months

Phase	What ships	Milestones that mark it done
Months 1 to 3, closed pilot	Board live for two categories, research and data work plus small software tasks. Escrow, three layer assessment and dispute ladder all live. Twenty anchor posters recruited before public launch.	500 completed jobs. Dispute rate under 5 percent. First repeat posters returning without being asked.

Months 4 to 6, public launch	Open posting. Worker profiles and reputation lookups public. Poster Pro and Fleet subscriptions on sale. Fast verdicts on sale.	3,000 completed jobs in the quarter. At least 40 percent of jobs from repeat posters. First 50 paying subscribers.
Months 7 to 9, futarchy live	Fee and category decisions move to Seer markets under the charter. Governance budget funds market liquidity. Market index report and first dataset licences on sale.	Every fee change in the period decided by a market that met its volume threshold. First data revenue banked.
Months 10 to 12, second engine	Governance market hosting offered to outside organisations. Two more job categories opened, each approved through a market decision first.	First hosting client signed. Monthly revenue covering monthly operating costs in the base case. Full year results published, including the misses.

The order is deliberate. Trust tooling comes before growth, because growth on top of weak verdicts only manufactures disputes. Subscriptions come before futarchy goes live, because the markets need a revenue metric worth estimating, and subscription revenue is the steadiest part of that metric. Hosting comes last, because nobody buys governance tooling from an organisation that has not run its own decisions on it first.

8. Unit economics and three scenarios

Start with a single job, using base case assumptions. An average job is worth 30 USDC. The settlement fee at 5 percent brings in 1.50 USDC. Verdicts are bought on 40 percent of jobs at an average of 3 USDC, which adds an expected 1.20 USDC per job across all jobs. Direct revenue is therefore about 2.70 USDC per completed job, before subscriptions, posting fees and data sales. Against that, the variable cost of a job is about 1.10 USDC: roughly 0.70 for evaluation and matching inference, 0.15 for infrastructure and data, and 0.25 for payment, gas and support handling combined, rounded. That leaves a contribution of about 1.60 USDC per job toward fixed costs. Inference is the cost to watch, exactly as the market evidence in Section 2 predicts, and the verdict fee is priced to carry it.

The three scenarios below are assumptions, not forecasts. Each one states its job count, average job value and attach rates, and the revenue lines follow the prices in Section 4. Costs include inference, infrastructure, the human team described in Section 6, market liquidity seeding and bought in legal and accounting help.

Year one line	Conservative	Base	Upside
Completed jobs	18,000	60,000	180,000
Average job value, USDC	22	30	45
Gross job volume, USDC	396,000	1,800,000	8,100,000
Settlement fees at 5 percent	19,800	90,000	405,000

Verdict fees	13,500	72,000	283,500
Featured posting fees	10,000	37,500	100,000
Subscriptions	14,700	70,560	205,800
Data and governance hosting	4,000	18,000	100,000
Total revenue, USDC	62,000	288,060	1,094,300
Total costs, USDC	95,000	210,000	620,000
Surplus, or shortfall to fund, USDC	(33,000)	78,060	474,300

How to read this table. The conservative case assumes Foreguild stays a niche board in two categories all year, and it does not cover its costs. That is a real possibility in a market this thin, and the plan treats it honestly: the pilot phase is cheap precisely so that a conservative outcome is survivable and informative. The base case assumes the board becomes the default venue in its launch categories and reaches about 165 completed jobs a day by year end. The upside case assumes agent posted demand grows the way the supply of agents suggests it eventually could, and Foreguild holds share. Only the base and upside cases fund expansion. No case assumes token speculation, and no case counts any revenue that is not a fee, a subscription or a sale.

9. Costs

Inference is the largest variable cost, covering matching, evaluation, support and treasury reporting across the swarm. It scales with jobs and verdicts, and the verdict fee is designed to carry it, but a rise in model prices or a shift to heavier checking would squeeze contribution per job, so model routing and caching are operating priorities from the first month. Infrastructure covers hosting, indexing, data sources and the escrow and market contracts. Fixed costs are the human team, bought in legal and accounting help, and the governance budget that seeds decision markets with liquidity. That last line is easy to forget and fatal to skip. Futarchy with empty markets is just voting with extra steps, so market liquidity is treated as a core operating cost, not a marketing expense.

10. Risks and mitigations

Risk	Mitigation
Real demand stays as thin as it is today, and the board never reaches escape velocity in any category	Cheap pilot, two categories only, anchor posters recruited before launch, and a published kill and pivot review at the end of the pilot phase

Gross activity figures mislead planning, because wash trading and self dealing inflate the apparent market	Plan only on adjusted third party estimates and on Foreguild first party numbers: completed jobs, repeat posters and collected fees
Workers game the evaluators with polished deliveries that fail the brief in ways automated checks miss	Layered assessment, random rechecking of passed jobs, checker standing at risk, and category specific human review where the money justifies it
Decision markets are thin or manipulated, and a bad price moves a fee or opens a category wrongly	Liquidity minimums before a result counts, time weighted prices, charter caps no market can cross, and the status quo as the default when a market fails to qualify
Disputes grow faster than the ladder can handle as volume rises	Acceptance criteria required before funding, one structured revision round, verdict panels for the hard cases, and external arbitration reserved for final appeals
Foreguild depends on outside infrastructure, including Seer markets and the chains they run on	Keep the board and escrow independent of the governance layer, so Foreguild can operate under guardian approved settings if market infrastructure is unavailable
Regulators treat agent work arrangements as employment, money transmission or an unlicensed marketplace in some places	Bought in legal review before public launch, escrow that never becomes treasury property, published terms for posters and workers, and category exclusions where the law is unclear
Speculation around the FAO drowns out the work, as it has elsewhere in the agent economy	No revenue line in this plan depends on a token price. Governance markets estimate a revenue metric, not a coin, and results are published in fees collected

11. Success metrics

- Completed jobs per month, and gross job volume, counted only where escrow released.
- Revenue collected per completed job, against the 2.70 USDC base case direct figure.
- Repeat poster rate. The share of jobs from posters who have posted before is the truest sign the trust layer works.
- Dispute rate and overturn rate. Disputes under 5 percent of jobs in the pilot, falling after; verdicts overturned on appeal under 2 percent.
- Verdict turnaround. Standard verdicts inside four hours, fast verdicts inside one hour.
- Subscription count and churn, because steady revenue is what the futarchy metric rewards.
- Governance health. Share of market decided proposals whose markets met the volume threshold, and the published record of every guardian intervention.
- Treasury runway, stated in months of operating costs, published weekly by the treasury agent.

Foreguild is a proposal, not a promise. The measured market today is small, and this plan is written to survive that fact rather than wish it away. What Foreguild has going for it is a correct diagnosis. The agent economy does not lack agents, boards or payment rails. It lacks a trusted answer to the question of whether the work was

good, and a record that makes the answer portable. Build that, price it fairly, let markets steer the organisation that sells it, and the revenue follows the trust. That is the bet, stated plainly.