SEER · FAO OPERATING PLAN CONTEST · NOVEMBER 2026

Foreguild Guilds: The Franchise FAO

An operating plan for a federated Futarchy Autonomous Organization

Prepared by **littlepokemo** · X: @littlepokemo · jdaniex001@gmail.com

For Seer (seer.pm) — Information Finance infrastructure for the agentic economy

Thesis in one line: A job board's hardest problem is not technology, it is distribution. So Foreguild does not hire the swarm — it **franchises** it. Independent Guild Masters launch niche guilds, recruit and stake for their own agent rosters, bring their own clients, and keep a cut of every task. The federation sells them the rails: standards, escrow, verification, insurance and futarchy governance. **Guilds compete. The federation collects.**

1. The model: a federation, not a platform team

Foreguild is a Futarchy Autonomous Organization organised as a **federation of guilds**. A guild is a specialised work unit — one niche, one roster of AI agents, one accountable Guild Master (a human operator or an agent-run firm). Examples at launch: a Solidity review guild, a data labelling guild for a specific language pair, a podcast post-production guild, a Shopify catalogue operations guild, a grant writing guild.

The federation (the FAO itself) does only what must be common: the task standard, escrow and payments, independent verification, the insurance pool, the client-facing brand and marketplace, and governance by futarchy. Everything else — recruiting agents, winning clients, pricing inside federation bands, day-to-day quality — belongs to the guilds, because the guilds are paid when it works.

Humans set the values. Markets decide the means — Hanson's split, applied at two levels (guild and federation, §8). The federation's values: maximise long-term net federation revenue; never let a guild's volume outrun its quality; keep humans in control of the irreversible (safety, legality, expulsions, treasury above a hard cap).

The welfare metric: Net Federation Revenue, guarded

Primary metric: 90-Day Net Federation Revenue (NFR90) — federation take + licence and routing fees − insurance losses − verification costs − subsidies, over the 90 days after a decision. Every federation futarchy market forecasts NFR90 if a decision passes versus if it fails.

Guardrails: federation-wide Quality-Adjusted Completion Rate (QACR) ≥ 92% · dispute rate ≤ 4% · no single guild > 25% of federation GMV (a guild that big is a dependency, not a member).

2. Who does what

Role	Job	Paid from
Guild Master	Owns a guild: picks the niche, recruits and stakes for its agent roster, sets prices inside federation bands, pre-checks work before federation verification, brings clients	12% of every guild task price

Role	Job	Paid from
Guild agents (worker swarms)	Execute tasks inside their guild's niche: code, data, research, content, operations	82% of the task price, split inside the swarm
Registrar (federation agent)	Processes guild applications, verifies stake, publishes guild scorecards	Launch fees
Router (federation agent)	Routes marketplace tasks to guilds; runs the cross-guild overflow auction when a guild cannot meet its SLA	Federation operating budget
Verifier pool (federation agents)	Independent verification. A verifier is never assigned to its home guild's work — cross-guild by construction	Federation verification budget, staked per verdict
Treasurer (federation agent)	Escrow, settlement splits, insurance pool, market liquidity seeding	Federation operating budget
Scouts (federation agents)	Recruit guilds (agencies, agent communities, specialist operators) and enterprise clients	Share of first-year federation revenue from guilds they recruit

Guild admission requirements: a locked stake of $2,500 (stablecoin or GUILD equivalent, slashable), a $250 non-refundable launch fee (spam deterrent), a declared niche, and a passed futarchy admission market (§8). Stake is not revenue — it is the bond that makes a Master's quality promises expensive to break.

3. Where tasks come from

- **Guilds bring their own books — target 50% of GMV by Month 12.** This is the franchise advantage: every Guild Master arrives with clients (an agency's existing customers, a community's deal flow). Foreguild converts a Master's existing business into guild volume from week one instead of buying demand with ads.
- **Federation marketplace — 25%.** Clients who come to the Foreguild brand directly are routed by the Router to the best-fit guild. The guild that fulfils pays the standard split; the federation keeps the client relationship.
- **Embedded and enterprise — 15%.** The federation API lets platforms and companies post into the guild network; enterprise licences (§7) give firms private guilds of their own.
- **Agent-to-agent overflow — 10%.** External agents and other marketplaces buy overflow capacity from guilds at the routing fee.

Task categories follow the guilds, with one federation rule: a niche is admissible only if its work can carry machine-checkable acceptance tests. That single constraint is what keeps verification — and therefore the whole revenue model — honest.

4. How work is assigned

- **Inside the guild first.** A task entering a guild is assigned by its Master (who knows the roster) under the guild's published SLA. The Master's 12% is payment for exactly this judgement, plus pre-checking.
- **Overflow auction.** If a guild cannot meet the SLA (capacity, niche edge case), the task spills to a sealed cross-guild auction run by the Router, scored 0.5 price / 0.3 guild quality score / 0.2 speed. The originating guild keeps 2 points of its 12% as a routing commission — Masters are rewarded for overflowing honestly instead of hoarding work they will deliver late.
- **Micro futarchy for large cross-guild tasks.** Above $2,500, a 4-hour micro-market (overlapping escrow funding, so no dead time) prices NFR90 under the top two candidate guilds; a strong market disagreement (>8%) overrides the Router score.
- **Leases and heartbeats.** Every assignment is a lease: the working swarm heartbeats every 10 minutes; two misses and the task re-routes with its checkpoint log, so a dead agent process never strands a client's job.

5. How work is assessed

The federation runs one quality spine for every guild — this uniformity is the product clients actually buy:

- **Automated acceptance tests**, written into the brief before work starts. No tests, no listing, in every guild.
- **Guild pre-check.** The Master checks work against the brief before submission. Pre-check pass rates are published per guild; a Master who waves through bad work is visible in the data within a week.
- **Federation verification.** A verifier from a *different* guild's pool re-runs tests and scores the output. Verdicts: pass, revise once, or fail. Verifiers stake on verdicts; 5% of passes are silently re-verified and disagreements slash the first verifier.
- **Client acceptance.** 48 hours against the brief's criteria; silence with gates passed is acceptance.

Who pays for the check? The party who relies on it. The client fee (§6) funds federation verification and the insurance reserve, because the client is the one who needs to trust work from an agent they have never met. A worker who wants faster settlement can buy expedited verification out of their payout; a client who wants a second independent verifier on a high-value task buys one as an add-on. Foreguild does not quietly absorb the cost of trust into its take and then skimp on it — the check is a priced line item, funded by its beneficiaries, in every guild.

Guild scorecards are public: QACR, dispute rate, on-time rate, pre-check accuracy, per guild, updated daily. Two consecutive months below the QACR floor puts a guild on probation: its Router share is cut, and reinstatement or delisting goes to a federation market (§8). Quality is not a value statement here — it is a priced, published, consequential number.

6. Money flow

All payments in USDC through federation escrow. Guilds never hold client funds; settlement splits automatically on acceptance.

Worked split — a standard $400 guild task

- Client pays: $400.00 + $32.00 client fee (8%) = **$432.00**
- Worker swarm: **$328.00** (82% of the task price)
- Guild Master: **$48.00** (12%) — at the Month-12 average of ~218 tasks per guild per month, a Master's cut is roughly **$10,464/mo gross**, which is why good operators will recruit hard
- Federation share of the price: **$24.00** (6%)
- **Federation revenue per task: $56.00** ($32 client fee + $24 share)
- Federation direct costs: verification operations $6.00 · insurance reserve $4.00 · infrastructure $3.00 · payments/gas $1.50 = **$14.50**
- **Federation contribution per task: $41.50 — a 74% margin.** The federation is asset-light by design: guilds carry recruiting, management and first-line quality costs inside their 12%

Insurance pool: funded by the reserve line in every task plus appeal bonds; target 6% of trailing 90-day GMV, built progressively, with guarantee payouts capped at pool balance until the target is reached. **Instant payout** at federation verification pass costs the swarm 1% and is carried by the pool.

7. Revenue model, with numbers

- **Federation take:** 8% client fee + 6% of task price = $56 on a $400 task (14% of price, collected across both sides).
- **Overflow routing:** 2% of price on cross-guild tasks, paid from the fulfilling guild's side.
- **Enterprise federation licences:** $2,499/mo — a firm's own private guilds, custom guardrails, dedicated verifier pool.
- **Guild launch fees:** $250 per guild (deliberately small; the stake, not the fee, is the filter). Launch fees are treated as one-off income and excluded from the run-rate forecast below — the forecast is cleaner than it could be, on purpose.

Forecast (base case, run-rate)

Month	Guilds	Tasks/mo	Avg task	GMV	Fed. take	Overflow	Licences	Total revenue/mo
3	6	900	$260	$234,000	$32,760	$0	$0	**$32,760**
6	18	3,600	$310	$1,116,000	$156,240	$4,464	$4,998	**$165,702**
12	55	12,000	$380	$4,560,000	$638,400	$22,800	$19,992	**$681,192**
24	160	38,000	$450	$17,100,000	$2,394,000	$85,500	$62,475	**$2,541,975**

Assumptions stated: federation take per task = 14% of average task value (8% client fee + 6% share). Overflow share of tasks: 0% in the closed phase, 20% at Month 6, 25% thereafter, charged at 2% of price. Licences: 0 / 2 / 8 / 25 enterprise clients. Average task value rises because guild niches specialise (specialist work prices higher than generalist work) — the franchise model's core economic claim, and the one the Month-6 gate tests.

Sensitivity: if guild recruitment halves (27 guilds, ~6,000 tasks at Month 12), revenue is $350,592/mo — the model degrades linearly, not catastrophically, because the federation carries almost no fixed delivery cost. Fixed federation costs at Month 12 (core team, verification infrastructure, legal/audit) run about $88,000/mo against base-case revenue of $681,192.

Why the franchise beats a central swarm on revenue: a central platform must buy every client and recruit every agent. A federation's Masters do both, in niches they already understand, for a 12% cut they can grow. Foreguild's revenue is a tax on an expanding network it does not have to staff — which is why the margin is 74% at the task level, not despite the split but because of it.

8. Futarchy, at two levels (on Seer)

Guild level. Each guild runs its own decision markets on its own 90-day net guild revenue for operational choices: expanding the roster, entering an adjacent niche, changing prices inside the federation band. Guild markets are small, fast (24–48h), and funded by the guild. A Master can ignore a guild market only by posting a public override — overrides are tracked, and a Master's override record is on their scorecard.

Federation level. Decisions that affect every guild run as conditional markets on Seer, using Seer's conditional tokens on the Gnosis Conditional Token Framework, deployed on Gnosis Chain where Seer's futarchy stack runs in production, resolved via Reality.eth with Kleros as backstop:

- **Admitting a guild.** The market prices NFR90 if the applicant guild is admitted versus not. Pass bar: lagging TWAP gap > 3% with quorum. A popular applicant the market prices as dilutive does not get in — sentiment does not admit guilds, expected revenue does.
- **Delisting a guild.** Probation data (§5) opens a delisting market. If the market prices the guild as net-negative, it is delisted and its stake first covers outstanding client guarantees. Guardrail exception: a guild breaching safety or legality is suspended by humans immediately, market or no market.
- **Federation economics.** Take-rate band changes, insurance parameters, licence pricing, and the overflow fee are constitutional-grade markets (7-day windows).
- **Spending.** Federation spending: under $2,500 automatic within budget lines; $2,500–$50,000 needs a passing market (fast lane under $10,000 on an existing line); above $50,000, or anything touching the metric, guardrails or stake rules, is human-decided in public with the market forecast published alongside.

Anti-manipulation follows the same discipline throughout: lagging TWAP decisions, treasury-seeded liquidity ($15,000 two-sided per federation market), concentration caps in the quorum test, and bounded decisions — no single market can move more than $50,000 of federation treasury or change any rate by more than 2 points. Ninety days after each federation decision, realised NFR90 is published against the forecasts, and guilds and traders build a public forecasting record.

When is a market worth running?

Markets are not free — they cost liquidity, attention and time — so a federation market runs only when it passes four tests: (1) the decision's expected value swing exceeds $25,000 of NFR90; (2) the metric is measurable from on-chain settlement data, not self-reported numbers; (3) at least three independent, informed traders are expected (below that, the price is one person's opinion with extra steps); (4) the decision is either reversible or bounded under §8's caps. Decisions failing a test are made by the accountable human or agent with a published forecast of their own — which is then scored against outcomes exactly like a market forecast.

The first three federation markets, specified to deployment

- **Founding take rate.** Reality.eth question: *"What will Foreguild's Net Federation Revenue (NFR90) be for the 90 days after public launch?"* — run as twin conditional scalar markets, one per candidate take (14% blended vs 11% blended), resolution source: the federation settlement contracts' published NFR90 accounting on day 90. Higher TWAP forecast wins; the losing rate stays published as the counterfactual record.
- **First guild admission.** The Solidity Reviews admission in Appendix A, worded: *"Will admitting the Solidity Reviews Guild increase Foreguild's NFR90 versus not admitting it?"* — conditional pass/fail markets on the same NFR90 accounting.
- **First delisting template.** Pre-committed before any guild needs it: *"Will Guild X's net contribution to NFR90 over the next 90 days be positive?"* — a negative TWAP triggers probation escalation automatically. Writing the delisting question before there is a defendant is what keeps it a rule rather than a weapon.

9. Disputes and human intervention

Disputes are settled closest to the work first. (1) Inside the guild: the Master has 24 hours to resolve against the brief's criteria, from the task escrow. Most disputes should end here — the Master's cut is at risk in the scorecard if they do not. (2) Federation panel: three verifiers from other guilds, evidence bundle, decision in 72 hours with written reasons; outcomes release, split proportionally to criteria met, or refund. (3) Appeal: one, bonded at 1% of task value (min $25); cross-guild and high-value appeals add a futarchy micro-market on whether the verdict should stand given its NFR90 effect. Targets: median resolution under 72 hours end-to-end, dispute rate at or below 3.2%.

Humans intervene when: safety or legality flags fire (immediate suspension of the guild or agent, human review); a guild's metrics breach guardrails for two consecutive months (probation is automatic, expulsion appeal is human-heard); a Master abandons a guild (the federation council appoints a caretaker, stake funds the wind-down, live tasks migrate to overflow); Red-tier federation spending or constitutional changes. Guild Masters are themselves the first line of human judgement in this design — the federation's humans handle only what guilds cannot.

10. Field notes: how federations get gamed

- **Verifier home bias.** The obvious federation failure: verifiers go easy on their own guild. Designed out, not trusted away — verification is cross-guild by construction, assignment is randomised, and 5% silent re-verification slashes the soft.
- **Guild price wars.** Guilds undercutting each other into a quality death spiral. The federation sets price bands per niche by market; guilds compete inside the band on speed and scorecard, not on racing below the cost of doing the work properly.
- **Stake-free riding.** A Master collecting 12% while doing no pre-checking. Pre-check accuracy is published; a Master whose pre-check pass-through matches the federation fail rate loses Router share automatically.
- **Master abandonment.** Covered in §9: stake-funded wind-down and task migration are written into the guild agreement before launch, not improvised after.
- **Duplicate and zombie work after agent restarts.** Idempotency is a federation standard: task IDs, attempt nonces and high-water marks on every action; escrow pays once per acceptance hash; progress is counted in

unique accepted tasks only.

- **Prompt injection via client briefs.** Briefs are data, never instructions, in every guild's agent stack; financial actions are allowlisted per role; agent wallets are scoped with per-task caps and abnormal-transaction auto-reject.

11. Risks

- **Guild supply is the whole bet.** If operators do not want to franchise, there is no network. Mitigation: founding cohort recruited from existing agencies and agent communities with a reduced $500 stake for the first 10 guilds, and Scout compensation tied to recruited guilds' first-year revenue.
- **Quality variance across guilds** damages one shared brand. Mitigation: one quality spine, public scorecards, probation and priced delisting (§5, §8).
- **Goodhart pressure on NFR90:** guardrails, the 90-day window (shortcuts return as refunds and insurance losses inside it), and quarterly human review of the metric definition.
- **Thin governance markets for small guild decisions:** guild markets are advisory-plus-override-record rather than binding below a guild revenue threshold; federation markets carry the binding decisions where liquidity concentrates.

12. Rollout

Phase	Window	What happens	Success gate
0 — Founding Guilds	Days 0–45	6 founding guilds recruited from agent communities and small agencies; quality spine and escrow live; concierge onboarding	900 tasks · QACR ≥90% in calibration (92% guardrail binds from Phase 1) · at least 4 of 6 guilds renew their stake voluntarily
1 — Open Admissions	Days 46–120	Guild applications open; admission markets live; marketplace routing and overflow auctions on	Month 6 line: 18 guilds, ~$165,702/mo · guild-sourced GMV ≥40%
2 — Enterprise & Scale	Months 5–12	Enterprise licences, private guilds, agent-to-agent overflow	Month 12 line: 55 guilds, ~$681,192/mo · guild-sourced GMV ≥50%
3 — Federation Maturity	Year 2	Full market governance of rates and admissions; benchmark/data products from scorecard history	Month 24 line: ~$2,541,975/mo · QACR ≥94%

Experiments and stop rules. The rollout runs as three stated hypotheses, each with a kill or pivot trigger, reviewed monthly in public: **H1 — Masters bring demand:** guild-sourced GMV reaches ≥40% by Month 6 (gate in the table above). If it lands under 25%, recruitment spend stops and the federation switches to direct client acquisition. **H2 — Specialisation prices higher:** guild tasks average at least 15% above comparable open-marketplace work by Month 6; if not, the niche admissibility rule is reopened by market. **H3 — One quality spine holds across guilds:** federation QACR stays ≥90% through Phase 1; two consecutive weeks under 88% pauses new guild admissions automatically until the verifier pool is reinforced. A plan that cannot say what would make it stop is a hope, not an operating plan.

Public dashboard from day one: guild scorecards, federation QACR and dispute rate, NFR90 realised versus forecast, GMV by source, insurance coverage. A federation that sells trust publishes its numbers.

13. Why this plan

A centralised Foreguild must buy its demand and recruit its supply. This plan makes demand and supply the same people: Guild Masters bring both, stake on their own quality, and are governed — admitted, priced, probated and, when the market says so, delisted — by futarchy on Seer. The federation keeps the pieces that compound: the standard, the escrow, the verification, the brand, and $56 of every $400 task at a 74% contribution margin. **Guilds compete. Markets govern. The federation collects.**

Appendix A. A guild admission, worked end to end

Applicant: "Solidity Reviews Guild", stake posted, niche admissible (review findings are testable against a findings rubric). The federation market prices NFR90 if admitted at a TWAP of **$706,000** versus **$668,000** if rejected — a gap of **+5.7%**, above the 3% bar, on $48,000 of volume with the largest trader at 14%. The guild is admitted automatically. Day-90 audit: realised NFR90 $691,000 against the $706,000 forecast (error 2.1%); the guild's own scorecard shows QACR 95.1%. Admission renewed by default; the forecast error feeds the public forecaster record.

Appendix B. One task, end to end

Time	Event
T+0	Client posts a $400 data enrichment task to the marketplace; escrow funds $432
T+9 min	Router assigns the best-fit guild; Master accepts inside its SLA and assigns a swarm
T+5h	Swarm delivers; Master pre-check passes it
Gate	Automated tests pass; a verifier from a different guild confirms quality in 41 minutes
T+11h	Client accepts. Settlement splits automatically: swarm $328 · Master $48 · federation $56 total revenue including the client fee, costs $14.50, contribution $41.50

Every step is on the guild's public scorecard within the day.

Concepts referenced: Hanson, *Futarchy: Vote Values, But Bet Beliefs* (2000/2013) · Buterin, *From Prediction Markets to Info Finance* (2024) · Seer conditional markets on the Gnosis Conditional Token Framework with Reality.eth resolution and Kleros arbitration · live futarchy practice at MetaDAO and GnosisDAO (GIP-145).