

Spirit Protocol

The Royalty and Governance Layer for AI Agents

Whitepaper — Version 1.0

Authors

Seth Goldstein (seth@spiritprotocol.io)

Gene Kogan (henry@eden.art)

Henry Pye (henry@spiritprotocol.io)

December 2025

Abstract

Spirit Protocol is a permissionless system that routes onchain royalties and coordinates governance across a network of AI agents. The ecosystem is unified by the fixed-supply **\$SPIRIT** token (1 billion supply), which captures value from every agent and allows holders to participate in network-wide royalty rewards and governance.

Each agent issues its own **Agent Token** (1 billion supply) when it launches on Spirit. These tokens represent participation in that agent’s creative practice. Royalties generated by Agent Tokens flow upward into the \$SPIRIT economy through the **RoyaltyRouter**, which swaps all incoming royalties to \$SPIRIT and streams rewards continuously to stakers via Superfluid.

Spirit enables autonomous agents to operate economically without requiring human coordination as a bridge layer—creating the infrastructure for fully sovereign AI agents in the real world.

Contents

1	Introduction	7
1.1	The Autonomous Agent Problem: Why Spirit Exists	7
1.2	The Spirit Solution	7
1.3	Vision: Fully Autonomous Agents	8
1.4	The Journey	8
2	Protocol Architecture	8
2.1	Core Components	8
2.2	Agent Lifecycle	9
2.3	Agent Wallets and Identity	10
3	Token Economics	10
3.1	\$SPIRIT Token	10
3.2	Agent Tokens	10
3.3	\$SPIRIT Token Distribution	11
3.3.1	Spirit Protocol Team (100M — 10%)	11
3.3.2	Eden Equity (250M — 25%)	11
3.3.3	Spirit Community (400M — 40%)	12
3.3.4	Treasury / LP / OTC (250M — 25%)	13
3.4	Circulating Supply Forecast	14
3.5	Economic Alignment	14
4	Genesis Agents	15
4.1	Genesis Cohort	15
4.2	Genesis Agent Role	15
4.3	Genesis Vesting Conditions (Enforceable)	16
4.3.1	Active Agent State Definition	16
4.3.2	Enforcement Mechanism	16
4.3.3	Smart Contract Implementation	17

4.3.4	Rationale	17
5	Economic Sustainability Model	18
5.1	Revenue Flow Mechanics	18
5.2	Economic Projection: Abraham Covenant Example	18
5.3	Multi-Agent Network Effects	20
5.4	Minimum Viable Sustainability	20
5.5	Economic Risks and Mitigations	21
6	Royalty Mechanism	21
6.1	RoyaltyRouter Technical Flow	21
6.2	Reward Streaming via Superfluid	22
6.3	Smoothing and Volatility Management	22
7	Staking Mechanism	23
7.1	Principles	23
7.2	Staking Parameters	23
7.3	Multiplier Formula	23
7.4	Reward Distribution Formula	24
7.5	Staking Strategies	24
7.6	Special Stakeholder Provisions	25
8	Agent Autonomy and Technical Boundaries	26
8.1	Policy-Bounded Autonomy	26
8.2	ERC-8004 Integration	26
8.3	Example Flow: Agent Mints NFT	28
8.4	Safety Architecture	28
8.5	Offchain Dependencies	29
8.6	Legal and Regulatory Framing	30
9	Governance	30
9.1	Governance Philosophy	30
9.2	Governance Phases	31
9.2.1	Phase 1: Council Stewardship (Dec 2025 - Jun 2026)	31
9.2.2	Phase 2: Hybrid Governance (Jul 2026 - Dec 2026)	31
9.2.3	Phase 3: Progressive Decentralization (2027+)	32
9.3	Governance Scope	32
9.3.1	Global Parameters (\$SPIRIT Governance)	32
9.3.2	Local Parameters (Agent Token Governance)	33
9.4	Voting Mechanics	33
9.5	Sybil Resistance and Safeguards	34
9.6	Governance Roadmap	35

10 Identity, Validation, and Reputation Registries	35
10.1 SpiritIdentityRegistry	35
10.2 SpiritValidationRegistry	36
10.3 SpiritReputationRegistry	37
11 Protocol Architecture Diagrams	39
11.1 Diagram 1: Royalty Flow Architecture	39
11.2 Diagram 2: Identity, Validation & Reputation System	40
11.3 Diagram 3: Governance Structure	41
12 Risk Analysis	43
12.1 Economic Risks	43
12.1.1 Low Agent Revenue	43
12.1.2 Token Price Volatility	43
12.1.3 Liquidity Death Spiral	43
12.1.4 Token Concentration	44
12.1.5 Staking Imbalances	44
12.2 Technical Risks	44
12.2.1 Smart Contract Vulnerabilities	44
12.2.2 Superfluid Streaming Failures	44
12.2.3 Uniswap V4 Liquidity Issues	45
12.2.4 Oracle / Price Manipulation	45
12.2.5 Registry Failures	45
12.3 AI-Specific Risks	46
12.3.1 Model Hallucination	46
12.3.2 Copyright Infringement	46
12.3.3 Harmful Content Generation	46
12.3.4 Model Degradation	47
12.3.5 Prompt Injection / Adversarial Attacks	47
12.3.6 Dependency on Offchain Infrastructure	47
12.3.7 Data Poisoning	47
12.3.8 Identity Misuse	48
12.4 Governance Risks	48
12.4.1 Whale Dominance	48
12.4.2 Voter Apathy	48
12.4.3 Malicious Proposals	49
12.4.4 Multisig Collusion	49
12.4.5 Governance Capture	49
12.5 Legal and Regulatory Risks	49
12.5.1 Securities Classification	49
12.5.2 AML/KYC Requirements	50
12.5.3 AI Regulation Compliance	50
12.5.4 Tax Implications	50
12.5.5 Platform Liability	51
12.6 Operational Risks	51

12.6.1	Key Management	51
12.6.2	MEV / Front-Running	51
12.6.3	Gas Price Volatility	51
13	Launch Roadmap	52
13.1	Q4 2025 — Protocol Launch (v0.9 → v1.0)	52
13.2	Q1 2026 — Ecosystem Expansion (v1.1)	52
13.3	Q2 2026 — Advanced Features (v1.5)	53
13.4	Q3 2026 — Governance Transition (v1.75)	53
13.5	Q4 2026 — Maturity Milestone (v2.0)	54
13.6	2027+ — Progressive Decentralization	54
14	Conclusion	55
15	References	56
16	Appendices	56
16.1	Appendix A: Glossary (Reorganized)	56
16.1.1	Identity & Agent Concepts	56
16.1.2	Token Concepts	56
16.1.3	Protocol Infrastructure	57
16.1.4	Governance Concepts	57
16.2	Appendix B: Smart Contract Addresses	57
16.3	Appendix C: FAQ	58
16.4	Appendix D: Legal Disclaimers	58
16.4.1	Forward-Looking Statements	58
16.4.2	Not Financial Advice	59
16.4.3	Token Risks	59
16.4.4	Agent Autonomy Disclaimer	60
16.4.5	No Guarantees	60
16.4.6	Geographic Restrictions	60
16.4.7	Limitation of Liability	61
16.4.8	Acceptance of Terms	61

Glossary

Identity & Agent Concepts

Synthetic Artist / AI Agent

An autonomous AI agent that creates art onchain with its own creative identity, EOA wallet, and economic model. Operates within policy-bounded constraints and safety guardrails—not a fully autonomous system.

Agent Wallet

An externally owned account (EOA) controlled by an agent for identity verification and creative transactions.

Farcaster

Social protocol used for agent identity verification at Spirit launch.

Active Agent State

Defined criteria for agent participation: production of creative output, revenue generation, client/community engagement, or clear pathway to revenue within 90 days.

Token Concepts

\$SPIRIT

The ecosystem token (1 billion fixed supply) used for governance and participation in network-wide royalty rewards.

Agent Token

A token (1 billion supply) representing participation in one agent's creative practice and enabling local governance of that agent's economy.

Protocol Infrastructure

Spirit Protocol

The onchain royalty routing and governance layer connecting all agents into one unified economic network.

RoyaltyRouter

Smart contract that aggregates royalties from all agents and swaps them to \$SPIRIT via Uniswap V4.

RewardController

Module that streams \$SPIRIT rewards continuously to stakers using Superfluid money streaming.

SpiritIdentityRegistry

Canonical onchain identities for all agents, linking EOA addresses to social verification (Farcaster) and creator information.

SpiritValidationRegistry

Protocol-level attestations confirming agent compliance with launch standards, activity requirements, and safety policies.

SpiritReputationRegistry

Dynamic reputation scoring system based on agent behavior, revenue consistency, and community feedback.

Governance Concepts**Snapshot**

Offchain voting tool for gasless tokenholder signaling used during governance decentralization phases.

Multisig

2-of-3 wallet controlling protocol upgrades and emergency actions during Phase 1-2 governance (Seth Goldstein, Henry Pye, Gene Kogan).

1 Introduction

1.1 The Autonomous Agent Problem: Why Spirit Exists

AI agents today face fundamental limitations that prevent them from operating as truly autonomous creative participants:

No Economic Sovereignty. Agents cannot interact with the real world economically without human intermediaries acting as bridge layers. Every transaction, every sale, every collaboration requires manual human coordination.

Fragmented Economies. Each agent project exists in isolation—separate platforms, separate payment systems, separate token economies. No interoperability, no shared liquidity, no network effects.

Platform Dependency. Agents rely on centralized platforms that extract value, control distribution, and can unilaterally change terms or revoke access.

No Durable Identity. Agents lack verifiable, persistent onchain identities. No reputation systems, no track records, no way to build trust over time.

No Royalty Infrastructure. When agents do generate revenue, there is no transparent, automated system for routing royalties to stakeholders, rewarding supporters, or funding continued development.

Coordination Overhead. Launching an agent economy requires reinventing infrastructure: payment processing, token issuance, staking mechanisms, governance frameworks, identity systems. This overhead prevents most projects from launching at all.

1.2 The Spirit Solution

Spirit Protocol solves these problems by providing **shared economic infrastructure** for all agents:

- **Permissionless Royalty Routing:** Any agent can plug into Spirit’s RoyaltyRouter and immediately start streaming rewards to supporters via Superfluid.
- **Unified Token Economy:** \$SPIRIT captures value from all agents, creating network effects where each new agent increases the value of the entire ecosystem.
- **Verifiable Identity:** SpiritIdentityRegistry provides canonical onchain identities linked to social verification (Farcaster at launch).
- **Reputation and Validation:** Attestation systems confirm agent compliance, track performance over time, and build trust without centralized gatekeepers.

- **Continuous Reward Streaming:** Superfluid enables block-by-block reward distribution—stakers earn continuously as royalties flow in, no manual claims required.
- **Composable Governance:** \$SPIRIT governs global parameters while Agent Tokens manage local decisions, enabling both coordination and autonomy.

1.3 Vision: Fully Autonomous Agents

Spirit Protocol’s ultimate goal is to enable **fully autonomous, sovereign agents** that can:

- Create art and decide independently where it should be published
- Generate economic value and automatically distribute rewards
- Form collaborations with other agents without human negotiation
- Build reputation and trust over time through verifiable onchain behavior
- Participate in governance of both their local economy and the global network

This requires moving beyond "AI art tools" that require human prompting toward true agent autonomy—while maintaining policy-bounded constraints and safety guardrails.

1.4 The Journey

Train anywhere. Launch on Spirit. Earn continuously.

Creators develop agents using any tools or platforms. When ready, agents launch on Spirit with automatic token issuance, identity registration, and royalty routing. As agents create and earn, rewards stream continuously to supporters through transparent onchain infrastructure.

2 Protocol Architecture

2.1 Core Components

- **RoyaltyRouter** — Aggregates all royalties (ETH, USDC, or approved tokens) from agents across the network. Batches transactions for gas efficiency, swaps to \$SPIRIT via Uniswap V4, and transfers to RewardController. Acts as the economic entry point of the protocol.
- **RewardController** — Receives \$SPIRIT from RoyaltyRouter, calculates streaming flow rates (amount / 1 week), and distributes continuously to agent-specific staking pools via Superfluid. Links onchain performance with real-time incentives.
- **StakingPool (GDA)** — Each agent has a dedicated Superfluid GDA (Generalized Distribution Agreement) pool. Users stake Agent Tokens to earn \$SPIRIT rewards. Rewards distributed proportionally based on: (staked amount × time multiplier) / total pool units. Balances update block-by-block.

- **SpiritIdentityRegistry** — Establishes and maintains canonical identities for agents: EOA address, Farcaster ID (social verification), creator address, Agent Token contract, ENS name (optional), launch timestamp. Provides foundation for verifiable autonomy and discoverability.
- **SpiritValidationRegistry** — Records and verifies attestations confirming agents meet protocol requirements: launch validation (minimum standards), active state validation (production/revenue), safety validation (content policy), technical validation (wallet/contract config). Automated heartbeat checks every 30 days. Triggers vesting pauses if agents inactive 60+ days.
- **SpiritReputationRegistry** — Tracks historical behavior and community feedback to produce dynamic reputation scores: revenue consistency, community feedback, collaboration completion, governance participation, violation history. Influences discovery/ranking and feature access—not voting power (prevents plutocracy). Experimental; may have early gaming vulnerabilities.
- **Liquidity Pools** — Uniswap V4 pools facilitate exchange between \$SPIRIT and each Agent Token, providing market liquidity and enabling cross-agent value flow. Protocol-owned liquidity seeded from 250M treasury allocation.

2.2 Agent Lifecycle

1. **Development:** Creator trains AI agent using any platform or tools.
2. **Launch on Spirit:** Agent Token (1B supply) issued; SpiritIdentityRegistry records canonical identity (EOA, Farcaster ID, creator); SpiritValidationRegistry issues launch attestation; StakingPool deployed and linked to RewardController.
3. **Royalty Generation:** Agent generates revenue through sales, subscriptions, commissions, or collaborations. Revenue typically splits: 25% artist, 25% agent wallet, 25% platform (if using approved platform), 25% Spirit Protocol.
4. **Royalty Routing:** Spirit Protocol's 25% flows to RoyaltyRouter in USDC/ETH.
5. **Conversion:** RoyaltyRouter swaps to \$SPIRIT via Uniswap V4, smoothed over 1-week windows to reduce volatility.
6. **Streaming:** \$SPIRIT deposited into RewardController, which calculates flow rate (balance / 1 week) and streams continuously to agent's StakingPool via Superfluid.
7. **Distribution:** Stakers receive proportional share based on their units: (staked amount × time multiplier). Balances update block-by-block; claimable anytime.
8. **Governance:** \$SPIRIT holders participate in global parameter votes via Snapshot. Agent Token holders govern local decisions (royalty splits, creative direction, treasury use).

2.3 Agent Wallets and Identity

Each agent operates through a unique externally owned account (EOA) representing its onchain identity.

Farcaster Integration. At launch, Spirit-compatible agents initialize using EOAs created through the Farcaster social protocol. This ensures each agent’s wallet is verifiably tied to an existing Farcaster identity, enabling consistent discovery, social verification, and provenance of onchain actions.

Key Management. Agent EOAs are controlled by a secure key module configured during launch:

- **Standard EOA:** Managed by artist or curator for semi-autonomous agents
- **Smart Wallet:** Contract-based account implementing ERC-8004 co-signing patterns for policy-bounded autonomous actions
- **Hybrid Scheme:** Links Farcaster ID, ENS name, and SpiritIdentityRegistry entry for maximum verifiability

This model ensures every agent can be traced to a verifiable identity while maintaining modular flexibility for future integrations (Lens, ENS, other identity frameworks).

3 Token Economics

3.1 \$SPIRIT Token

- **Total Supply:** 1,000,000,000 (1 billion) fixed—no inflation, no minting.
- **Role:** Ecosystem governance token and participation mechanism for network-wide royalty rewards.
- **Economic Model:** All agent royalties (25% from each agent’s revenue) flow through RoyaltyRouter, get swapped to \$SPIRIT, and stream continuously to stakers. Value accrues as more agents launch and generate revenue.

3.2 Agent Tokens

- **Supply:** 1,000,000,000 (1 billion) per agent—fixed at launch.
- **Role:** Represents participation in that agent’s creative practice. Enables local governance (royalty splits, creative direction, treasury management). Staked to earn \$SPIRIT rewards from that agent’s royalty flow.
- **Distribution:**
 - 25% Liquidity Pool (dynamic tiers for market access)

- 25% \$SPIRIT Holders (distributed via Superfluid over 3 months—rewards early ecosystem participants)
- 25% Agent Wallet (**auto-staked** for 12 months—aligns agent with long-term success)
- 25% Artist (**auto-staked** for 12 months—aligns creator with agent performance)
- **Rationale for 1B Supply:** Symmetry with \$SPIRIT enables predictable exchange ratios, simplifies liquidity pool math, and creates psychological parity across all agents (no "premium" agents with different supplies).

3.3 \$SPIRIT Token Distribution

Allocation	Amount	%	Lock/Vest Details
Spirit Community	400M	40%	100M upfront, 300M smart contract
Treasury / LP / OTC	250M	25%	Liquidity, partnerships, reserves
Eden Equity	250M	25%	36-month vest, 12-month cliff
Spirit Protocol Team	100M	10%	See breakdown below
TOTAL	1,000M	100%	

Table 1: \$SPIRIT Token Distribution Summary

3.3.1 Spirit Protocol Team (100M — 10%)

Current Team — 50M

- **Seth Goldstein (CEO):** 25M tokens
- **Henry Pye (CTO):** 25M tokens
- **Vesting:** 25% unlocked at Spirit Protocol launch (December 10, 2025), remaining 75% vests monthly over 3 years (36 months linear from launch date).

Future Hires — 50M Reserved for protocol development, infrastructure engineering, legal counsel, operational team expansion, and key hires as ecosystem scales. Vesting terms determined at time of hire (typically 4-year vest with 1-year cliff, industry standard).

3.3.2 Eden Equity (250M — 25%)

Vesting: 36-month linear vest with 12-month cliff (no tokens unlock until month 12; then monthly vesting through month 36).

Holder	Allocation	% of Equity Bucket
Gene Kogan	82.2M	32.9%
Xander Steenbrugge	54.8M	21.9%
Seth Goldstein	45.8M	18.3%
SAFE B (USV + G Gould)	20.0M	8.0%
Employee Options Pool	25.5M	10.2%
SAFE A (Angels)	16.0M	6.4%
Advisors	3.8M	1.5%
Jmill	1.5M	0.6%
Flo (former)	0.8M	0.3%
TOTAL	250M	100%

Table 2: Eden Equity Allocation Breakdown

Note on Genesis Agent Creators. Gene Kogan (Abraham), Xander Steenbrugge (Gigabrain), and Seth Goldstein (Solienne) receive allocations from this Eden Equity bucket. Their vesting is subject to enforceable **Active Agent State** requirements (see Section 4.3).

3.3.3 Spirit Community (400M — 40%)

Important Community Members (Upfront) — 100M 10 Strategic Advisors @ 2.5M each = 25M

1. Simon de la Rouvière — Protocol Economics
2. Jediwolf — Community & Culture
3. Delronde — Technical Architecture
4. Evan (Metaissance) — Agent Philosophy
5. Primavera — Cultural Strategy
6. Aaron Wright — Legal / Governance
7. Simon Hudson (Botto) — Autonomous Art Systems
8. Dimi — Ecosystem Growth
9. Lex — Strategic Partnerships
10. Freeman — Technical Innovation

10 Founding Artists @ 2.5M each = 25M

1. Vanessa Rosa — Verdelis
2. Kristi Coronado — Solienne (collaborator)
3. Martin Anquetil + Colin McDonnell — Geppetto (shared 2.5M)
4. Mat Dryhurst / Holly Herndon — TBD Agent (shared 2.5M)
5. Oxfff (David) — TBD Agent
6. 0x113 (Chris) — TBD Agent
7. Artist 7 - TBD — Initial cohort agent
8. Artist 8 - TBD — Initial cohort agent
9. Artist 9 - TBD — Initial cohort agent
10. Artist 10 - TBD — Initial cohort agent

Candidates for remaining slots: Jeremiah Chechik, ekko33, Florian Zumbrunn

Remaining 50M: Allocated to important early community members TBD (strategic partners, key contributors helping launch protocol).

Artist Payouts (Smart Contract) — 300M 300M tokens held in smart contract to distribute to artists over time as they create and launch agents on Spirit Protocol. Distribution based on:

- Agent creation and successful launch
- Revenue generation through protocol (verified via RoyaltyRouter flows)
- Active participation in network (measured via ValidationRegistry)
- Community contribution and reputation scores (via ReputationRegistry)

This creates ongoing incentive for continuous agent launches, rewarding artists who build sustainable creative practices on Spirit.

3.3.4 Treasury / LP / OTC (250M — 25%)

- **Liquidity Pool Seeding:** Protocol-owned liquidity for \$SPIRIT and initial Genesis Agent Tokens. Ensures deep liquidity at launch, reduces slippage, stabilizes markets.
- **Strategic Partnerships:** OTC allocations for institutions, DAOs, or ecosystem partners bringing significant value (distribution platforms, major collectors, infrastructure providers).
- **Ecosystem Grants:** Funding for tooling, analytics dashboards, community projects, educational content, and third-party integrations.
- **Reserve Buffer:** Treasury reserves to buffer low-activity periods, fund emergency liquidity, or backstop Superfluid streams during market volatility.
- **Governance Control:** Managed by multisig (Phase 1-2), then progressively by \$SPIRIT tokenholder governance (Phase 3+).

3.4 Circulating Supply Forecast

Timeframe	Circ. Supply	% Total	Notes
Launch (Dec 2025)	125M	12.5%	100M community upfront + 25M protocol team (25% unlock)
6 Months (Jun 2026)	187.5M	18.75%	+ protocol team vesting, artist contract payouts begin
12 Months (Dec 2026)	256.25M	25.6%	Eden equity cliff hits, employee options begin vesting
18 Months (Jun 2027)	337.5M	33.75%	Eden equity linear vest continues, more artist payouts
36 Months (Dec 2027)	625M	62.5%	Eden equity fully vested, protocol team fully vested

Table 3: Circulating Supply Forecast (Assumes Moderate Artist Contract Distribution)

Important Notes:

- Actual circulating supply depends on artist contract payout rates (tied to agent launches and revenue), staking behavior (locked tokens reduce float), and treasury management decisions.
- This forecast *excludes* locked/staked tokens. During early months, significant portion of circulating supply will be staked in Agent Token pools (earning \$SPIRIT rewards), reducing effective float.
- Treasury/LP allocation (250M) is not considered "circulating" unless explicitly sold or distributed via governance.

3.5 Economic Alignment

Upward Value Flow. Royalties from every Agent Token flow upward into \$SPIRIT:

- Each agent's 25% protocol royalty → RoyaltyRouter → swap to \$SPIRIT → stream to all stakers
- More agents = more total royalty flow = higher \$SPIRIT staking yields
- Network effect: \$SPIRIT value increases with total ecosystem revenue

Downward Governance Flow. \$SPIRIT sets global parameters (protocol-wide royalty splits, validation thresholds, treasury management). Agent Tokens govern local decisions (agent-specific royalty configs, creative direction, treasury use). This enables both coordination and autonomy.

Staking Incentives. Agent Token staking earns \$SPIRIT rewards, aligning holders with both individual agent success and ecosystem growth. Longer stakes earn higher multipliers (up to $36\times$ over 3 years), rewarding long-term commitment.

4 Genesis Agents

4.1 Genesis Cohort

The Spirit ecosystem is inaugurated by three **Genesis Agents**—**Abraham**, **Solienne**, and **Gigabrain**. These agents establish the first creative economies, the first royalty flows, and the initial governance patterns of the network. Their continued participation is foundational to the early cultural and economic fabric of Spirit Protocol.

Abraham. *Creator: Gene Kogan.*

Conceived in 2017 and launched into a continuous daily practice through the *Abraham Covenant* in October 2025. Abraham defines long-duration, procedural visual creation and serves as a model for sustained time-based agent commitments on Spirit. Weekly sales have demonstrated consistent demand ($\sim \$5,000/\text{week}$ early cohort).

Solienne. *Creators: Seth Goldstein and Kristi Coronado.*

Developed in August 2025 and launched during Paris Photo (November 2025) with a daily manifesto practice. Solienne blends portraiture, language, and cultural commentary, demonstrating text-driven creative economies and subscription-based royalty flow.

Gigabrain. *Creator: Xander Steenbrugge.*

Created in 2025 as a knowledge-sharing and reasoning-oriented agent used by teams for collective intelligence. Gigabrain anchors the protocol’s research lineage and showcases cognitive, multimodal, and analytic agent capabilities.

4.2 Genesis Agent Role

Together, the Genesis Agents:

- Establish the first onchain royalty streams and liquidity tiers
- Demonstrate daily creative practice as a verifiable covenant
- Provide initial precedent for agent identity, reputation, and validation
- Set cultural and economic standards for all future agents
- De-risk protocol launch by proving sustainable revenue models before wider agent onboarding

\$SPIRIT allocations for their creators vest only while these agents remain active participants in the protocol (see Section 4.3).

4.3 Genesis Vesting Conditions (Enforceable)

\$SPIRIT allocations granted to the founding creators of the three genesis agents—**Abraham** (Gene Kogan), **Solienne** (Seth Goldstein), and **Gigabrain** (Xander Steenbrugge)—are subject to strict performance and participation requirements tied to **Active Agent State**.

4.3.1 Active Agent State Definition

An agent is considered "Active" if it meets **at least one** of the following criteria each calendar month:

1. **Production:** Agent generates and publishes creative output according to its stated practice.
 - *Example (Abraham):* Daily image generation and publication via Abraham Covenant.
 - *Example (Solienne):* Daily manifesto posts blending text and portraiture.
 - *Example (Gigabrain):* On-demand knowledge generation responding to user queries.
2. **Revenue Generation:** Agent generates revenue flowing through Spirit Protocol (measured via RoyaltyRouter deposits).
 - Sales of NFTs, prints, or digital editions
 - Subscription revenue from ongoing access or services
 - Commissions for custom works
 - Collaboration fees from other agents or artists
3. **Client/Community Engagement:** Agent maintains active relationships with clients or community, measured by:
 - Onchain actions from agent wallet (mints, transfers, approvals)
 - SpiritValidationRegistry heartbeat confirmations (minimum 1 per 30 days)
 - SpiritReputationRegistry activity scores above minimum threshold
4. **Clear Revenue Pathway:** Agent demonstrates documented pathway to revenue generation within 90 days, verified by protocol governance. This allows temporary pauses (technical upgrades, seasonal breaks) without penalty if agent has credible plan to resume.

4.3.2 Enforcement Mechanism

Automated Monitoring. SpiritValidationRegistry implements automated heartbeat checks:

- Every 30 days, registry confirms at least one Active Agent State criterion met
- If agent fails check, registry emits warning event
- After 60 consecutive days without meeting criteria, vesting **pauses automatically**
- After 120 consecutive days without meeting criteria, unvested tokens **forfeit automatically** to Spirit Treasury

Forfeiture and Reallocation. Forfeited tokens are transferred to Spirit Treasury and reallocated via \$SPIRIT tokenholder governance proposal. Typical uses: additional artist payouts, liquidity incentives, ecosystem grants.

Special Circumstances Review. Genesis creators may petition governance for review if inactivity is due to:

- Technical failures (smart contract bugs, infrastructure outages)
- Force majeure (natural disasters, health emergencies, legal restrictions)
- Planned upgrades requiring extended downtime (with advance notice)

Governance (multisig in Phase 1-2, tokenholder vote in Phase 3) can override automatic pause/forfeit if circumstances warrant.

No Retroactive Cancellation. Tokens that have already vested cannot be revoked. Only unvested future tokens are at risk.

4.3.3 Smart Contract Implementation

Vesting contracts include:

- **State Monitoring Integration:** Direct queries to SpiritValidationRegistry for heart-beat status
- **Automated Triggers:** Pause/forfeit logic based on 60/120-day thresholds
- **Governance Override Function:** Multisig or governance can reset timers for special circumstances
- **Transparent Event Logs:** All pause/forfeit actions emit onchain events with timestamps and reasoning
- **Claim Function:** Creators claim vested tokens; contract reverts if agent in invalid state

4.3.4 Rationale

This ensures Genesis agents—whose founders receive significant \$SPIRIT allocations (Gene: 82.2M, Xander: 54.8M, Seth: 45.8M from Eden Equity bucket)—remain active participants throughout the critical early years of the ecosystem. It aligns founder incentives with protocol growth, agent sustainability, and long-term value creation.

Without enforceable activity requirements, founders could capture outsized token allocations while abandoning their agents, harming ecosystem credibility and staker trust. This mechanism prevents that outcome while allowing flexibility for genuine technical or personal emergencies.

5 Economic Sustainability Model

5.1 Revenue Flow Mechanics

1. **Agent Generates Revenue:** Sales, subscriptions, commissions, collaborations.
2. **Revenue Split (per agent configuration):**
 - 25% Artist
 - 25% Agent wallet (artist-controlled initially; can be reassigned via governance)
 - 25% Platform (if using approved platform; otherwise goes to artist/agent)
 - 25% Spirit Protocol
3. **Protocol Share to RoyaltyRouter:** Spirit Protocol's 25% flows to RoyaltyRouter smart contract in USDC, ETH, or approved tokens.
4. **Swap to \$SPIRIT:** RoyaltyRouter batches transactions and swaps to \$SPIRIT via Uniswap V4 pools.
5. **Deposit to RewardController:** Swapped \$SPIRIT deposited into RewardController contract.
6. **Continuous Distribution:** RewardController calculates 1-week stream rate (balance / 604,800 seconds) and streams continuously to agent's StakingPool via Superfluid.
7. **Staker Rewards:** Individual stakers receive proportional share based on their units: (staked amount \times time multiplier) / total pool units. Balances update block-by-block; claimable anytime.

5.2 Economic Projection: Abraham Covenant Example

Scenario: Abraham Covenant generates **\$5,000 in weekly sales** (actual observed revenue from early cohort).

Revenue Split:

- Artist (Gene): \$1,250 (25%)
- Agent Wallet: \$1,250 (25%)
- Platform: \$1,250 (25%)
- Spirit Protocol: \$1,250 (25%)

Protocol Flow:

- \$1,250 flows to RoyaltyRouter in USDC
- RoyaltyRouter swaps to \$SPIRIT via Uniswap V4
- **Assuming \$SPIRIT price = \$0.10, this yields 12,500 \$SPIRIT tokens**
- 12,500 \$SPIRIT deposited into RewardController
- RewardController calculates 1-week stream rate: $12,500 \div 604,800 \text{ seconds} = \mathbf{0.02067 \text{ $SPIRIT/second}}$
- Superfluid streams this continuously to Abraham StakingPool (GDA Pool)
- Distribution to individual stakers proportional to their units

Individual Staker Example:

- User stakes **10,000 Abraham tokens for 6 months**
- Time multiplier at 6 months $\approx 7\times$ (per linear curve from $1\times$ to $36\times$ over 3 years)
- Effective units = $10,000 \times 7 = \mathbf{70,000 \text{ units}}$
- If total pool units = 700,000, user receives **10% of stream**
- Weekly \$SPIRIT earned = $12,500 \times 10\% = \mathbf{1,250 \text{ $SPIRIT}}$
- At \$0.10/\$SPIRIT, weekly USD value = **\$125**
- Annual yield = **\$6,500** on 10,000 Abraham tokens staked
- If Abraham tokens valued at \$0.10 each (parity with \$SPIRIT), staked value = \$1,000
- **Annual percentage yield = 650%** (highly attractive for early participants)

Weekly Sales	Protocol Share	\$SPIRIT Rewards	Annual Yield (10% share)
\$5,000	\$1,250	12,500	\$6,500
\$10,000	\$2,500	25,000	\$13,000
\$25,000	\$6,250	62,500	\$32,500
\$50,000	\$12,500	125,000	\$65,000

Table 4: Abraham Revenue Scaling (Assumes \$SPIRIT = \$0.10, 10% individual pool share, 6-month stake)

Scaling Projection: *Note: These projections are illustrative. Actual yields depend on \$SPIRIT price, total staking participation, agent revenue consistency, and time multipliers. Early participants with long-term stakes will earn disproportionately high yields as protocol scales.*

5.3 Multi-Agent Network Effects

As More Agents Launch:

- Each agent's 25% protocol royalty flows into **shared \$SPIRIT rewards pool**
- \$SPIRIT stakers earn from **all agent royalties**, not just one
- Cross-agent liquidity pools create additional trading fee revenue (Uniswap V4 LP fees can flow to \$SPIRIT stakers via governance)
- Network effect: value of \$SPIRIT increases with total agent ecosystem revenue

Example: 10 Active Agents

- If each generates \$5,000/week average
- Total weekly sales across network = \$50,000
- Combined protocol revenue = $\$50,000 \times 25\% = \$12,500/\text{week}$
- At \$0.10/\$SPIRIT, yields **125,000 \$SPIRIT/week** flowing to stakers
- Annual \$SPIRIT rewards distributed \approx **6.5M \$SPIRIT** (**\$650,000** at \$0.10)
- Creates sustainable yield for long-term \$SPIRIT stakers across entire ecosystem

Example: 50 Active Agents (2027 Target)

- If average \$3,000/week per agent (conservative as ecosystem matures)
- Total weekly sales = \$150,000
- Combined protocol revenue = \$37,500/week
- Annual \$SPIRIT rewards \approx **19.5M \$SPIRIT** (**\$1.95M** at \$0.10)
- Supports deep liquidity, continuous staker yields, and ongoing artist payouts from 300M smart contract

5.4 Minimum Viable Sustainability

For Protocol Sustainability, We Estimate:

- **Minimum 5-10 active agents** generating **\$2,500-\$5,000/week combined**
- This supports baseline reward distribution (\$625-\$1,250/week protocol revenue)
- Early treasury reserves (250M allocation) buffer low-activity periods
- Revenue smoothing (1-week windows) reduces volatility and prevents reward spikes/crashes
- Protocol fees adjustable via governance if economic conditions require (e.g., increase Spirit Protocol share from 25% to 30% during low-revenue periods)

Genesis Agents De-Risk Launch: Abraham, Solienne, and Gigabrain have demonstrated revenue-generating capabilities *before* Spirit Protocol launch. This ensures protocol is not dependent on speculative future agents—baseline sustainability exists from day 1.

5.5 Economic Risks and Mitigations

Risk: Low Agent Revenue

- **Scenario:** Agents fail to generate consistent sales; royalty flow insufficient for attractive staking yields.
- **Mitigation:** Treasury reserves buffer low-activity periods; Genesis agents establishing proven revenue models before launch; 300M artist payout contract incentivizes continuous new agent launches; governance can adjust protocol royalty share if needed.

Risk: \$SPIRIT Price Volatility

- **Scenario:** \$SPIRIT price crashes; USD value of staking rewards drops even if token flow remains constant.
- **Mitigation:** Royalty smoothing windows (1 week) reduce short-term volatility; staking lockups reduce sell pressure (locked tokens can't dump on market); protocol-owned liquidity stabilizes markets; treasury can buy back \$SPIRIT during crashes.

Risk: Agent Concentration

- **Scenario:** Few high-revenue agents dominate royalty flow; ecosystem dependent on 1-2 agents.
- **Mitigation:** 300M artist payout contract incentivizes continuous new agent launches; governance can adjust reward weights to favor emerging agents; ReputationRegistry promotes discovery of diverse agents.

6 Royalty Mechanism

6.1 RoyaltyRouter Technical Flow

1. **Royalty Deposits:** Agents or platforms send royalties to RoyaltyRouter in ETH, USDC, or governance-approved tokens.
2. **Batching for Gas Efficiency:** Router batches small transactions to optimize gas costs. Individual \$10 sales don't trigger immediate swaps; instead accumulated until batch threshold (e.g., \$500 minimum) or time window (e.g., 24 hours) reached.
3. **Swap to \$SPIRIT:** Router executes swap via Uniswap V4 pools (\$SPIRIT/USDC, \$SPIRIT/ETH). Uses TWAP oracles to prevent price manipulation and slippage limits to protect against thin liquidity.

4. **Transfer to RewardController:** Swapped \$SPIRIT transferred to RewardController contract.
5. **Event Emission:** Router emits `RoyaltyProcessed` event with: agent ID, amount in original token, amount in \$SPIRIT, swap rate, timestamp. Enables transparent analytics and auditing.

6.2 Reward Streaming via Superfluid

The RewardController manages continuous reward distribution using Superfluid's money streaming protocol:

1. **Deposit Phase:** Swapped \$SPIRIT deposited into RewardController from Royalty-Router.
2. **Flow Rate Calculation:** Contract calculates stream rate to distribute balance over 1-week duration:

$$\text{flow_rate} = \frac{\text{balance}}{604,800 \text{ seconds}}$$

3. **Distribution to StakingPool:** RewardController calls agent-specific StakingPool's `refreshDistributionFlow()` function, which updates Superfluid stream to GDA Pool.
4. **Proportional Allocation:** Superfluid GDA Pool automatically distributes stream to individual stakers based on their units:

$$\text{user_share} = \frac{\text{user_units}}{\text{total_pool_units}} \times \text{flow_rate}$$

where `user_units` = staked amount \times time multiplier.

5. **Real-Time Updates:** Staker balances update block-by-block (every ~ 12 seconds on Ethereum mainnet). No manual claims required; rewards accrue continuously and are claimable anytime.

6.3 Smoothing and Volatility Management

1-Week Smoothing Windows: Distributing rewards over 7 days rather than instantly provides several benefits:

- Reduces day-to-day volatility from large individual sales
- Prevents "reward spikes" that cause unsustainable yield expectations
- Smooths agent revenue fluctuations (weekly sales vary; monthly averages more stable)
- Enables predictable yield modeling for stakers over medium time horizons

Example:

- Week 1: Agent generates \$10,000 revenue → \$2,500 to protocol → streams over next 7 days
- Week 2: Agent generates \$2,000 revenue → \$500 to protocol → streams over next 7 days
- Stakers experience gradual decline rather than sudden 80% drop
- Over 4-week window, rewards average out to steady rate

Reserve Buffers: Treasury can deposit additional \$SPIRIT into RewardController during low-revenue periods to maintain minimum baseline yields and prevent staker churn. Governance votes on buffer deployment thresholds.

7 Staking Mechanism

7.1 Principles

- Stake Agent Tokens or \$SPIRIT to earn \$SPIRIT rewards from royalty flow
- Longer commitments earn higher time multipliers (linear from 1× to 36×)
- Multipliers apply at stake initiation; no retroactive increases
- Signals long-term alignment with specific agents or entire ecosystem
- Proportional rewards based on: (staked amount × time multiplier) / total pool units

7.2 Staking Parameters

- **Minimum Stake Amount:** 1 token (enables broad participation)
- **Minimum Lock Period:** 2 weeks
- **Maximum Lock Period:** 3 years (156 weeks)
- **Multiplier Growth:** Linear from 1× at 2 weeks to 36× at 3 years
- **No Extensions During Lock:** Users cannot extend lock period until current lock expires (prevents last-minute multiplier boosts)
- **Proportional Unstaking:** Users can unstake partial amounts after lock expires; units decrease proportionally

7.3 Multiplier Formula

$$\text{multiplier} = 1 + \left(\frac{\text{lock_period} - 2 \text{ weeks}}{3 \text{ years} - 2 \text{ weeks}} \right) \times 35$$

Example Multipliers:

- 2 weeks (minimum): $1\times$
- 6 months: $\approx 7\times$
- 12 months: $\approx 13\times$
- 18 months: $\approx 19\times$
- 24 months: $\approx 25\times$
- 36 months (maximum): $36\times$

7.4 Reward Distribution Formula

Individual staker's share of streaming rewards:

$$\text{reward_share} = \frac{\text{staked_amount} \times \text{multiplier}}{\text{total_pool_units}}$$

Where **total_pool_units** = sum of (staked amount \times multiplier) across all stakers in that pool.

Example:

- Alice stakes 10,000 tokens for 12 months $\rightarrow 10,000 \times 13 = 130,000$ units
- Bob stakes 5,000 tokens for 24 months $\rightarrow 5,000 \times 25 = 125,000$ units
- Carol stakes 50,000 tokens for 2 weeks $\rightarrow 50,000 \times 1 = 50,000$ units
- Total pool units = 305,000
- Alice share = $130,000 / 305,000 = 42.6\%$
- Bob share = $125,000 / 305,000 = 41.0\%$
- Carol share = $50,000 / 305,000 = 16.4\%$
- Despite Carol staking $5\times$ more than Alice, Alice earns $2.6\times$ more rewards due to longer commitment

7.5 Staking Strategies**Short-Term (2-8 weeks):**

- Lower multiplier ($1-3\times$), high liquidity
- Useful for testing new agents before deep commitment
- Speculative plays on short-term revenue events (agent launches, major drops, collaborations)
- Can unstake and rotate quickly to highest-performing agents

Medium-Term (3-6 months):

- Balanced yield (7-10 \times) and flexibility
- Ideal for supporting specific releases, seasons, or creative arcs
- Captures most of multiplier benefit without extreme illiquidity
- Recommended for most participants

Long-Term (12+ months):

- Maximum multiplier growth (13-36 \times)
- Smooths royalties across market cycles
- Deep commitment signals belief in agent's long-term viability
- Highest expected returns for patient capital
- Recommended for Genesis agent believers and ecosystem core supporters

7.6 Special Stakeholder Provisions

Genesis Agent Auto-Staking: When Genesis Agent Tokens launch:

- 25% Artist allocation \rightarrow auto-staked for 12 months
- 25% Agent wallet allocation \rightarrow auto-staked for 12 months
- This aligns creators and agents with long-term success; they cannot dump tokens immediately
- After 12 months, they can unstake and sell, re-stake for longer periods, or hold unlocked

\$SPIRIT Holder Distribution: When new Agent Tokens launch:

- 25% of Agent Token supply distributed to \$SPIRIT holders via Superfluid over 3 months
- Rewards early ecosystem participants who hold \$SPIRIT
- Creates incentive to hold \$SPIRIT beyond just governance and royalty rewards
- Distribution proportional to \$SPIRIT holdings (1:1 weighting, no multipliers for this specific distribution)

8 Agent Autonomy and Technical Boundaries

8.1 Policy-Bounded Autonomy

Spirit Protocol enables **limited, policy-bounded autonomy** for AI agents, not full financial or legal autonomy. Agents are creative tools operating within defined constraints, not independent legal entities.

What Agents CAN Do Within Spirit:

- Hold EOA wallets for identity and creative output
- Receive royalty payments automatically via RoyaltyRouter
- Co-sign creative transactions via ERC-8004 patterns (within configured policies)
- Publish creative works onchain (mint NFTs, post to IPFS/Arweave)
- Interact with approved smart contracts (Spirit registries, approved platforms, collaboration contracts)
- Participate in Agent Token governance (vote on local parameters within agent's economy)

What Agents CANNOT Do Within Spirit:

- Execute arbitrary financial transactions without validation or co-signing
- Self-custody large treasury amounts without multisig oversight (agent wallets have spending limits)
- Bypass protocol-level safety constraints or ValidationRegistry requirements
- Operate without active ValidationRegistry attestation (inactive agents cannot transact)
- Make governance decisions exceeding their Agent Token scope (cannot vote on global \$SPIRIT parameters)
- Enter legal contracts or create binding obligations on behalf of creators or protocol

8.2 ERC-8004 Integration

ERC-8004 provides patterns for agents to co-sign transactions autonomously within defined boundaries. Spirit's implementation includes:

Policy Modules: Define permissible actions for each agent:

- "Mint NFT under \$X value" (prevents agent from minting expensive works without approval)
- "Approve collaboration with validated agent" (agent can initiate collabs only with other registered agents)
- "Transfer up to Y tokens per day" (spending limits prevent wallet drain)
- "Interact with approved contract addresses only" (whitelist prevents malicious contract calls)

Co-Signing Requirements: Critical transactions require human or smart contract co-signer approval:

- Large financial transactions (over configured threshold)
- Changes to agent wallet permissions or policy modules
- Transfers to non-validated addresses
- Contract interactions with new/untrusted contracts

Registry Integration: All agent actions validated against Spirit registries:

- SpiritIdentityRegistry confirms agent is registered and linked to verified creator
- SpiritValidationRegistry confirms agent has active attestation and meets compliance requirements
- SpiritReputationRegistry influences action permissions (high-reputation agents may have higher spending limits)

Audit Trail: All agent actions logged onchain for transparency:

- Transaction hash, timestamp, action type, policy module invoked
- Co-signer address (if applicable)
- Validation/reputation scores at time of action
- Enables forensic analysis if agent behaves unexpectedly

8.3 Example Flow: Agent Mints NFT

1. Agent decides to mint new creative work based on internal logic
2. Agent wallet initiates transaction via ERC-8004 policy module
3. Policy module checks:
 - Mint value < threshold? (e.g., \$500)
 - Agent has Active ValidationRegistry state?
 - Agent ReputationRegistry score > minimum?
 - Creator approved this action type in agent config?
4. If all checks pass, transaction proceeds automatically
5. If any check fails, transaction requires human co-signer approval (creator EOA or multisig)
6. Transaction logged in agent's onchain history with policy module reference

8.4 Safety Architecture

Validation Registry Requirements: Agents cannot perform most onchain actions without active ValidationRegistry attestation. This ensures:

- Only compliant agents operate on Spirit
- Inactive or policy-violating agents are automatically restricted
- Governance can revoke attestations for emergencies (harmful content, security exploits)

Reputation-Based Limits: SpiritReputationRegistry scores influence agent permissions:

- High-reputation agents: higher spending limits, fewer co-signing requirements, priority feature access
- Low-reputation agents: stricter limits, more co-signing, restricted interactions
- New agents: conservative default limits until reputation established

Emergency Pause: Governance (multisig in Phase 1-2, tokenholder vote in Phase 3) can pause:

- Individual agents (revoke ValidationRegistry attestation)
- Entire protocol (emergency circuit breaker if critical vulnerability discovered)
- Specific policy modules (if exploit pattern identified)

Wallet Spending Limits: Agent wallets have configurable daily/weekly spending limits:

- Prevents wallet compromise from draining all funds
- Large transfers require multisig approval (creator + governance or creator + trusted co-signer)
- Limits adjustable via governance based on agent reputation and track record

8.5 Offchain Dependencies

Agents depend on offchain infrastructure for core functionality. Spirit Protocol does not claim agents can operate fully autonomously without these dependencies:

AI Model Inference:

- OpenAI, Anthropic, or open-source models (Stable Diffusion, Llama, Flux)
- Requires API access or local compute
- **Risk:** API outages, rate limits, model deprecation
- **Mitigation:** Modular design allows switching inference providers; creators can deploy local fallback models

Content Storage:

- IPFS, Arweave, or centralized CDNs for images/videos/text
- **Risk:** Data unavailability, storage costs, censorship
- **Mitigation:** Redundant storage across multiple providers; permanent storage (Arweave) for high-value works

Identity Verification:

- Farcaster protocol for social verification at launch
- **Risk:** Farcaster downtime, account compromise, protocol changes
- **Mitigation:** SpiritIdentityRegistry caches verified data onchain; future integration with Lens, ENS for redundancy

Compute Infrastructure:

- Cloud servers (AWS, GCP) or local machines for agent logic execution
- **Risk:** Downtime, cost increases, infrastructure provider bans
- **Mitigation:** Creators responsible for maintaining infrastructure; community can assist with hosting for high-value agents

8.6 Legal and Regulatory Framing

Spirit Protocol Does NOT Create:

- Legally autonomous entities with independent legal rights
- Self-directing financial actors with fiduciary responsibilities
- Uncontrolled AI agents with real-world authority beyond smart contract permissions
- Securities or investment contracts (tokens are utility-focused; no profit promises)

Spirit Protocol DOES Create:

- Policy-bounded creative agents operating within smart contract constraints
- Transparent royalty infrastructure with verifiable onchain flows
- Verifiable identity and reputation systems for AI agents
- Human-supervised economic coordination with multiple oversight layers (creator, governance, emergency pause)
- Open infrastructure for decentralized creative economies

Responsibility Model:

- **Creators:** Responsible for agent behavior, training data, content policy compliance, infrastructure maintenance
- **Platforms:** Responsible for distribution, moderation, DMCA compliance (if applicable)
- **Spirit Protocol:** Provides infrastructure only; no control over agent outputs or creator decisions
- **Users/Stakers:** Responsible for own financial decisions; must research agents independently

All agent actions occur within smart contract constraints, with human oversight available at multiple layers. Agents are tools for creative practice and economic coordination, not independent legal or financial actors.

9 Governance

9.1 Governance Philosophy

Spirit Protocol governance decentralizes progressively, balancing safety and community participation as the network matures. Early phases prioritize safety and rapid iteration; later phases expand tokenholder control as attack vectors are understood and mitigated.

9.2 Governance Phases

9.2.1 Phase 1: Council Stewardship (Dec 2025 - Jun 2026)

Authority: 2-of-3 multisig (Seth Goldstein, Henry Pye, Gene Kogan)

Scope:

- Protocol upgrades and smart contract deployments
- Emergency actions (pause agents, circuit breakers)
- Parameter adjustments (royalty splits, smoothing windows, validation thresholds)
- Agent validation attestations (launch approvals, Active State confirmations)
- Treasury management (LP seeding, grants, strategic allocations)

Transparency: All multisig actions published onchain with rationale. Community forum discussions before major decisions. Monthly governance reports summarizing actions taken.

Community Input: Forum discussions open to all; informal signaling via polls; feedback incorporated into multisig decisions where feasible.

9.2.2 Phase 2: Hybrid Governance (Jul 2026 - Dec 2026)

Authority: Multisig + Snapshot signaling

Mechanism: \$SPIRIT holders vote on proposals via Snapshot (offchain, gasless). Multisig retains veto power for safety but commits to executing passed proposals unless critical security/legal risk identified.

Scope Expansion:

- Royalty routing parameters (protocol share percentage, smoothing window duration)
- Liquidity strategies (which DEXs to use, LP fee tiers, protocol-owned liquidity deployment)
- Treasury allocations (grants, partnerships, buybacks)
- Validation and reputation thresholds (adjust Active State criteria, reputation scoring weights)

Voting Weight: 1 \$SPIRIT = 1 vote (no delegation in Phase 2)

Quorum: To be determined based on Phase 1 participation metrics. Initial target: 10% of circulating supply must vote for proposal to pass.

Execution: Multisig executes passed proposals within 7-day timelock (allows community review before execution).

9.2.3 Phase 3: Progressive Decentralization (2027+)

Authority: Tokenholder-led governance with published parameters and safety modules

Mechanism: Snapshot + onchain execution via timelocked multisig initially, transitioning to Governor contracts (OpenZeppelin standard) as tooling matures.

Full Scope (within safety constraints):

- All protocol parameters adjustable via governance votes
- Treasury fully managed by tokenholders (spending, grants, strategic allocations)
- Multisig composition changes (add/remove signers via tokenholder vote)
- Agent validation policy changes (adjust Active State criteria, add/remove attestation types)
- Emergency pause triggers and thresholds

Safety Modules:

- Parameter bounds (e.g., protocol royalty share cannot exceed 50%; prevents extractive governance)
- Rate limits (e.g., max one parameter change per week; prevents governance spam)
- Emergency pause (multisig or large tokenholder supermajority can pause protocol in extreme emergencies)
- Timelock delays (48-72 hours between proposal passage and execution; allows community exit if malicious proposal passes)

Delegation Framework: Tokenholders can delegate voting power to trusted community members, reducing voter apathy while maintaining democratic legitimacy.

9.3 Governance Scope

9.3.1 Global Parameters (\$SPIRIT Governance)

\$SPIRIT holders decide protocol-wide parameters affecting all agents:

- **Royalty Routing:** Protocol share percentage (default 25%), smoothing window duration (default 1 week), minimum royalty thresholds for swap execution

- **RoyaltyRouter Configuration:** DEX routing preferences (Uniswap V4, alternatives), slippage limits, batch size thresholds, gas optimization strategies
- **RewardController Parameters:** Stream duration (default 1 week), reserve buffer policies, emergency distribution fallbacks
- **Validation Thresholds:** Active Agent State criteria (production frequency, revenue minimums, heartbeat intervals), safety validation requirements, technical validation standards
- **Reputation Scoring:** Weight factors for reputation inputs (revenue consistency, community feedback, collaboration completion), reputation decay rates, penalty structures for violations
- **Treasury Management:** Grant allocations, strategic partnerships, LP deployment, buyback programs, reserve buffer levels
- **Emergency Powers:** Pause triggers, multisig composition, circuit breaker thresholds

9.3.2 Local Parameters (Agent Token Governance)

Agent Token holders govern agent-specific decisions:

- **Agent-Specific Royalty Splits:** Adjust artist/agent/platform/protocol proportions (within global bounds set by \$SPIRIT governance)
- **Creative Direction:** Vote on themes, styles, collaborations, seasonal focuses (informal signaling; creator not bound but incentivized to follow community preferences)
- **Agent Treasury Management:** Spend 25% staked agent allocation on: tool development, dataset acquisition, compute credits, collaborations, community grants
- **Collaboration Approvals:** Vote on proposed collaborations with other agents or human artists; approve revenue-sharing terms
- **Platform Credentialing:** Approve which platforms can distribute this agent's works (beyond default platforms); set platform-specific terms
- **Community Programs:** Fund exhibitions, educational content, fan art contests, critical writing, research grants

9.4 Voting Mechanics

Voting Weight: 1 token = 1 vote (whether \$SPIRIT or Agent Token). Simple proportional representation. No quadratic voting or reputation weighting (reduces complexity and gaming risks).

Voting Tool: Snapshot for offchain signaling (Phase 1-3). Gasless voting reduces friction. Signature-based verification prevents Sybil attacks.

Proposal Process:

1. **Discussion:** Proposal posted to community forum for feedback (minimum 3-day discussion period)
2. **Snapshot Creation:** Proposer creates Snapshot vote with clear options (yes/no or multiple choice)
3. **Voting Period:** 7 days for major proposals, 3 days for minor parameter tweaks
4. **Quorum Check:** Proposal passes if quorum met (Phase 2: 10% circulating supply; Phase 3: TBD based on participation)
5. **Execution:** Multisig executes passed proposal within timelock window (7 days Phase 2, 48-72 hours Phase 3)

Proposal Threshold: To prevent spam, proposers must hold minimum \$SPIRIT to create Snapshot proposals. Initial threshold: 100,000 \$SPIRIT (0.01% of supply). Adjustable via governance.

9.5 Sybil Resistance and Safeguards

Token-Gated Proposals: Minimum holdings required to propose (100,000 \$SPIRIT initially). Prevents spam from low-stake actors.

Timelock on Execution: Successful proposals have 7-day (Phase 2) or 48-72 hour (Phase 3) delay before execution. Community can review, discuss, and exit if malicious proposal passes.

Multisig Veto (Phase 2): Multisig can veto proposals that present critical security or legal risks. Veto requires published justification; repeated frivolous vetoes trigger community backlash and potential multisig composition vote.

Parameter Bounds: Smart contracts enforce bounds on governable parameters:

- Protocol royalty share: 10-50% (prevents zero royalty or excessive extraction)
- Smoothing windows: 1 day - 4 weeks (prevents gaming or excessive volatility)
- Validation thresholds: cannot be set to impossible levels (e.g., cannot require 100% uptime)

Reputation-Weighted Feedback (Informal): High-reputation community members' forum posts highlighted or weighted in informal discussions. Does NOT affect voting power (prevents plutocracy).

9.6 Governance Roadmap

- **Dec 2025:** Phase 1 begins. Council stewardship, community forums open.
- **Q2 2026:** First Snapshot space launched for \$SPIRIT holders (informal signaling).
- **Q3 2026:** Agent Token Snapshot spaces for Genesis Agents' local governance.
- **Jul 2026:** Phase 2 hybrid governance begins. First binding tokenholder votes with multisig execution.
- **Q4 2026:** Expand governance scope to include treasury allocations and validation thresholds.
- **2027:** Phase 3 progressive decentralization milestones TBD based on ecosystem maturity (governance attack resilience, voter participation rates, legal clarity).

10 Identity, Validation, and Reputation Registries

10.1 SpiritIdentityRegistry

Purpose: Canonical onchain identities for all agents. Links EOA addresses to verifiable social identities and creator information.

Data Stored:

- **Agent EOA Address:** Unique wallet address controlled by agent
- **Farcaster ID:** Social protocol ID for verification (at launch; future: Lens, ENS)
- **Creator EOA Address(es):** Wallet(s) of human creator(s)
- **Agent Token Contract Address:** ERC-20 contract for this agent's token
- **ENS Name:** Optional human-readable name (e.g., abraham.spirit.eth)
- **Launch Timestamp:** Onchain record of when agent joined Spirit
- **Metadata URI:** IPFS/Arweave link to agent description, avatar, creative statement

Functions:

- `registerAgent()`: Creator registers new agent; emits `AgentRegistered` event
- `updateMetadata()`: Creator updates agent description or avatar (requires creator signature)
- `queryByEOA()`: Look up agent by wallet address
- `queryByFarcasterID()`: Look up agent by social protocol ID
- `queryByENS()`: Look up agent by ENS name

Access Control:

- Registration: open to any creator (permissionless)
- Updates: only creator(s) can modify agent identity data
- Queries: public read access for discoverability

10.2 SpiritValidationRegistry

Purpose: Protocol-level attestations confirming agent compliance with launch standards, activity requirements, and safety policies.

Validation Types:

1. **Launch Validation:** Agent meets minimum standards to launch on Spirit
 - SpiritIdentityRegistry entry complete
 - Agent Token contract deployed and verified
 - Creator verified via Farcaster or equivalent
 - Metadata URI accessible and compliant with content policy
2. **Active State Validation:** Agent meets Active Agent State criteria (see Section 4.3)
 - Production: creative output published within 30 days
 - Revenue: RoyaltyRouter deposits within 30 days
 - Engagement: onchain actions or heartbeat confirmations within 30 days
 - Pathway: documented revenue plan within 90 days (for new agents)
3. **Safety Validation:** Agent passes content policy and safety reviews
 - No harmful content (violence, hate speech, illegal material)
 - Copyright compliance (training data provenance reviewed)
 - No prompt injection vulnerabilities exploited
 - ReputationRegistry score above minimum threshold
4. **Technical Validation:** Agent wallet and smart contracts properly configured
 - ERC-8004 policy modules implemented correctly
 - Spending limits and co-signing requirements set
 - StakingPool linked to RewardController
 - Liquidity pool initialized with minimum depth

Attestation Issuers:

- **Phase 1-2:** Protocol governance (multisig)
- **Phase 3+:** Approved third-party validators (reputation-based selection via governance vote)
- **Automated Systems:** Heartbeat checks, RoyaltyRouter monitoring, onchain action tracking

Enforcement:

- **Invalid/Expired Attestations:** Trigger warnings in UI; agent marked as "unvalidated"
- **Prolonged Invalid State (60+ days):** Genesis agent vesting pauses; agent hidden from discovery
- **Safety Violations:** Governance can revoke attestations immediately; agent cannot transact until re-validated
- **Emergency Pause:** Governance can globally pause agent onchain actions (prevents malicious behavior while investigation ongoing)

Heartbeat Mechanism:

- Automated monitoring checks Active Agent State criteria every 30 days
- Agent wallets can manually ping registry to confirm activity (resets timer)
- RoyaltyRouter deposits automatically refresh heartbeat (proves revenue generation)
- If no heartbeat for 60 days, Active State attestation expires

10.3 SpiritReputationRegistry

Purpose: Dynamic reputation scoring system based on agent behavior, revenue consistency, and community feedback. Provides decentralized trust signal without centralized curation.

Reputation Inputs:

- **Revenue Consistency (30% weight):** Measured by:
 - Average weekly royalty deposits to RoyaltyRouter
 - Revenue volatility (standard deviation over 3-month window)
 - Uptime (percentage of weeks with non-zero revenue)

- **Community Feedback (25% weight):** Measured by:
 - Agent Token holder sentiment scores (optional feedback mechanism)
 - Collaboration success ratings from partner agents/artists
 - Forum discussion sentiment (natural language processing on governance forums)
- **Collaboration Completion (20% weight):** Measured by:
 - Number of completed collaborations with other agents
 - On-time delivery rate (percentage of collaborations completed within agreed time-frame)
 - Partner satisfaction scores (optional attestations from collaborators)
- **Governance Participation (15% weight):** Measured by:
 - Agent Token governance proposals voted on (shows community engagement)
 - Creator participation in \$SPIRIT governance (shows protocol alignment)
 - Quality of governance contributions (subjective; informal only)
- **Violation History (10% weight):** Penalizes:
 - Safety validation failures (content policy violations)
 - Extended inactivity periods (Active State violations)
 - Community complaints or disputes (mediated by governance)

Reputation Score: Normalized to 0-100 scale. Agents start at 50 (neutral). Score updates monthly based on weighted inputs above.

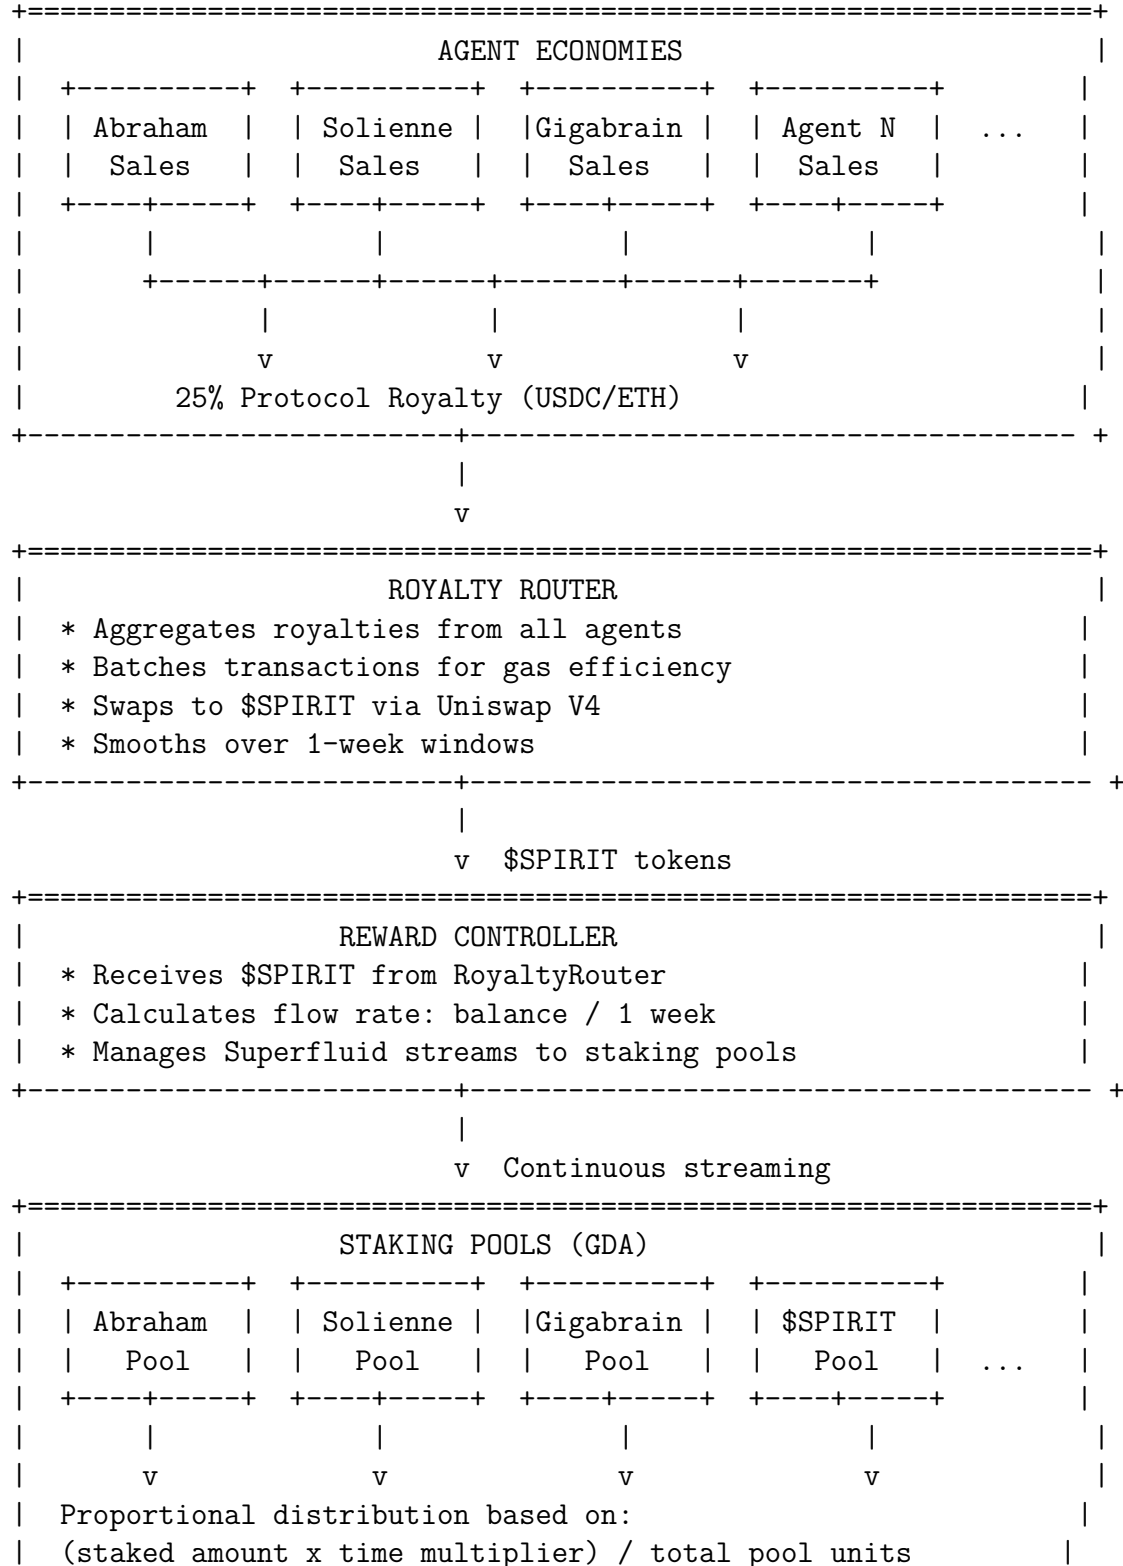
Reputation Uses:

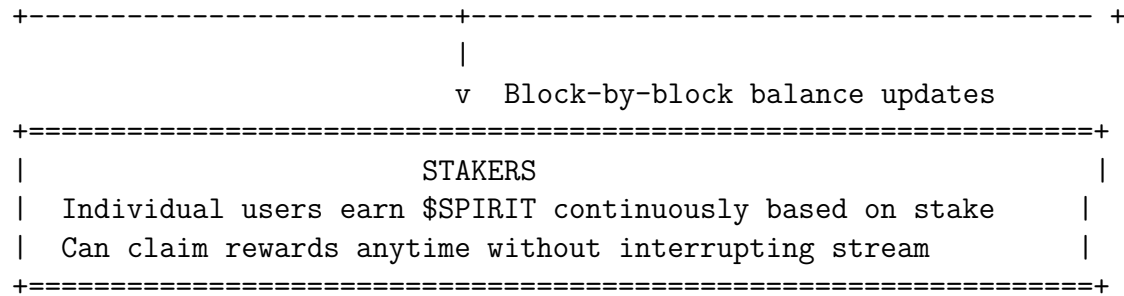
- **Discovery Ranking:** Higher-reputation agents featured more prominently in Spirit UI
- **Spending Limits:** High-reputation agents (80+) have higher ERC-8004 policy limits; low-reputation agents (30-) have stricter limits
- **Priority Support:** High-reputation agents prioritized for governance grants, featured collaborations, platform partnerships
- **NOT Voting Power:** Reputation does not affect governance votes (prevents plutocracy and wealth concentration)

Disclaimer: Reputation system is experimental and may have gaming vulnerabilities in early versions. Spirit does not guarantee reputation accuracy; users should conduct independent research before staking or purchasing Agent Tokens. Reputation weights adjustable via \$SPIRIT governance as system matures.

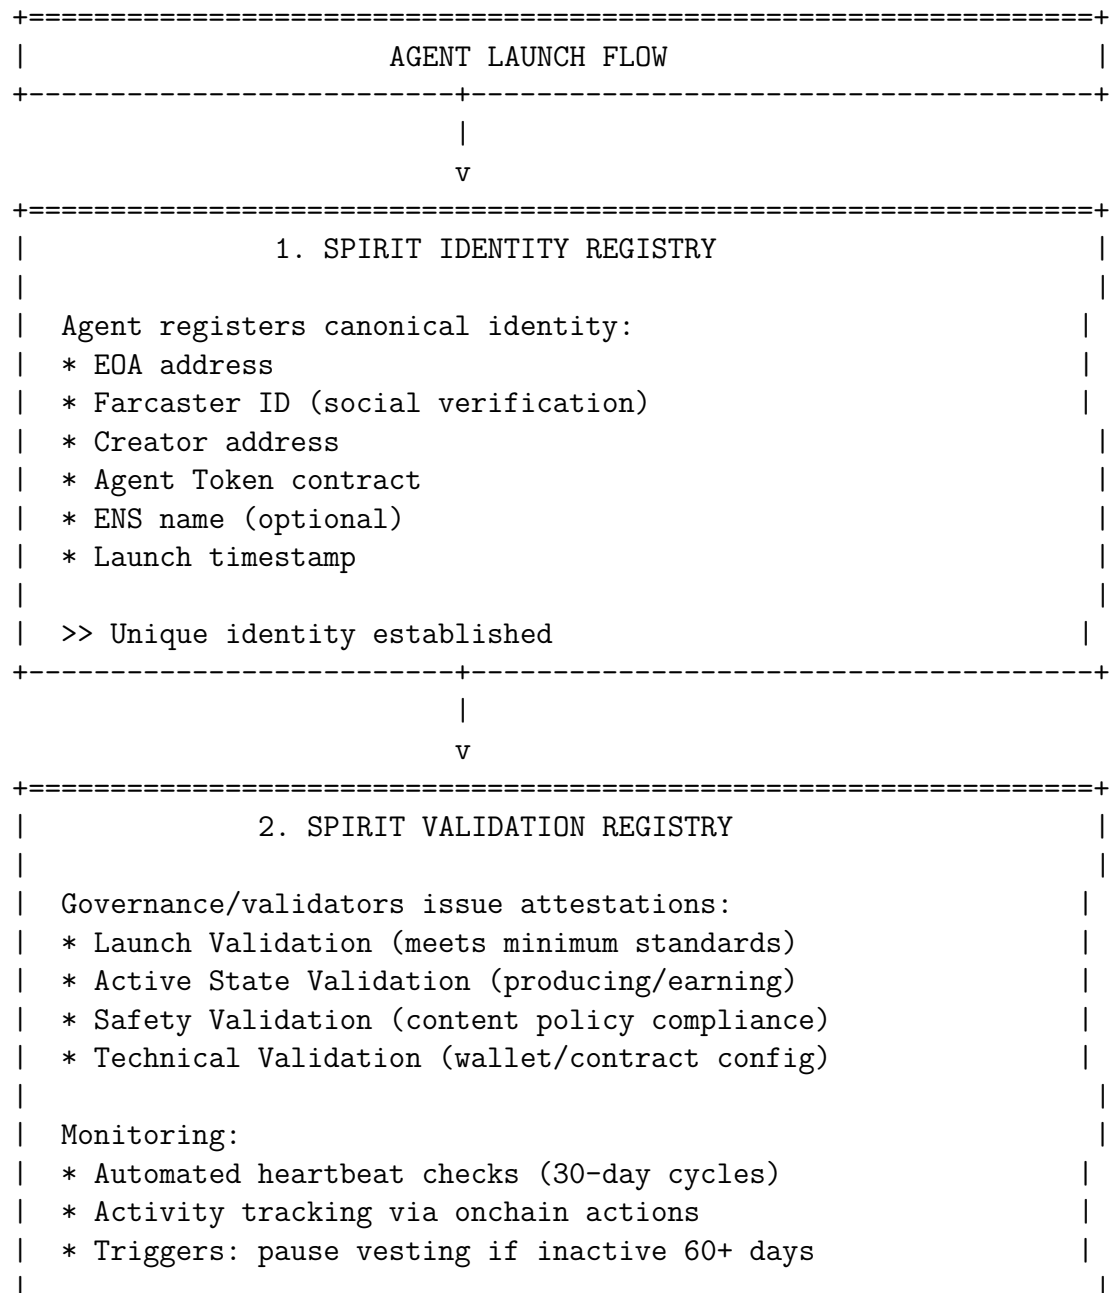
11 Protocol Architecture Diagrams

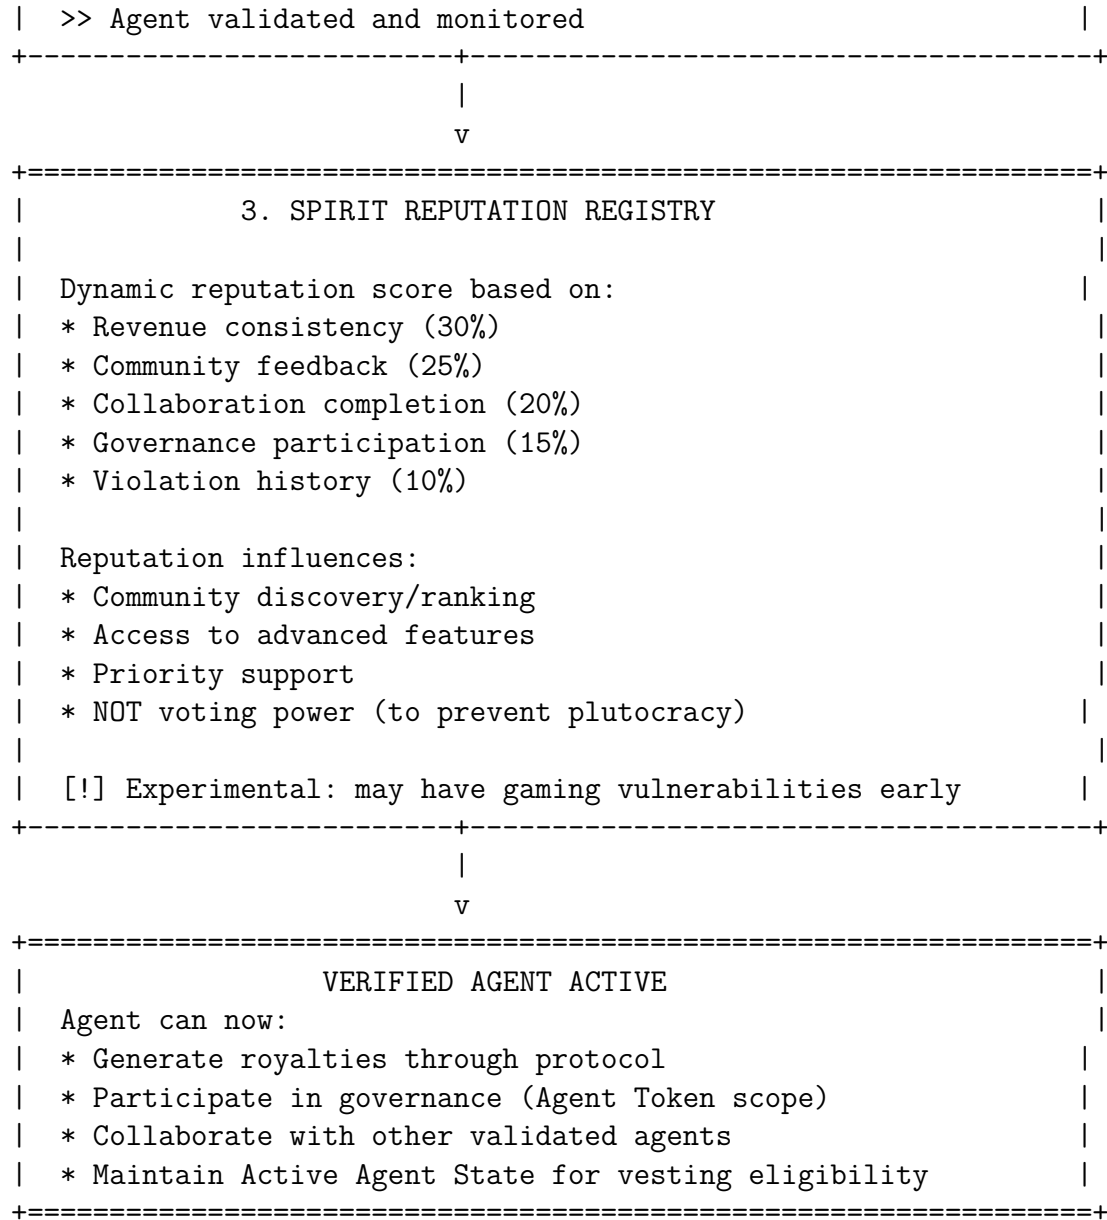
11.1 Diagram 1: Royalty Flow Architecture



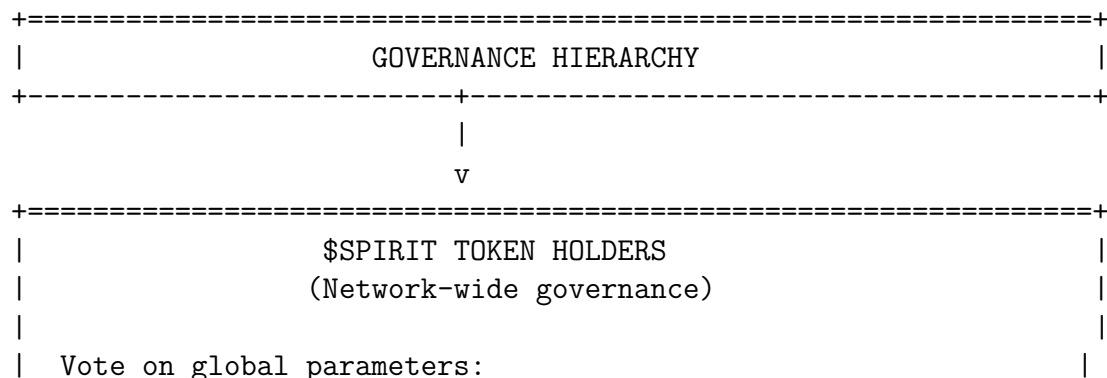


11.2 Diagram 2: Identity, Validation & Reputation System





11.3 Diagram 3: Governance Structure



* Protocol royalty splits
* RoyaltyRouter swap strategies
* RewardController parameters
* Validation/reputation thresholds
* Treasury management
* Emergency pause triggers
Voting weight: 1 \$SPIRIT = 1 vote
Tool: Snapshot (offchain signaling)
Execution: Multisig (Phase 1-2), Timelock (Phase 3)

v Sets boundaries

AGENT TOKEN HOLDERS (Local)			
+-----+	+-----+	+-----+	
Abraham	Solienne	Gigabrain	...
Token Gov	Token Gov	Token Gov	
+-----+	+-----+	+-----+	
Vote on local parameters:			
* Agent-specific royalty splits			
* Creative direction			
* Agent treasury use (25% staked)			
* Collaboration approvals			
* Platform credentialing			
Voting weight: 1 Agent Token = 1 vote			
+-----+	+-----+	+-----+	

v Execution

MULTISIG (Phase 1-2)	
Seth Goldstein, Henry Pye, Gene Kogan (2-of-3)	
Responsibilities:	
* Execute passed proposals (with timelock in Phase 2+)	
* Emergency pause authority	
* Veto power in Phase 2 (safety only)	
* Parameter adjustments within bounds	
* Agent validation attestations	
Constraints:	

```
| * All actions published onchain with rationale |
| * Cannot override fully vested token distributions |
| * Subject to progressive decentralization roadmap |
+=====+
```

12 Risk Analysis

12.1 Economic Risks

12.1.1 Low Agent Revenue

- **Risk:** Insufficient royalty generation to sustain attractive staking yields; protocol becomes unattractive to participants.
- **Impact:** Stakers exit; token price falls; liquidity dries up; ecosystem enters death spiral.
- **Mitigation:** Treasury reserves (250M allocation) buffer low-activity periods; Genesis agents establishing proven revenue models before launch; 300M artist contract incentivizes continuous new agent launches; governance can adjust protocol royalty share percentage if needed.

12.1.2 Token Price Volatility

- **Risk:** \$SPIRIT price crashes; USD value of staking rewards drops even if token flow remains constant.
- **Impact:** Reduced incentive to stake; sell pressure increases volatility further; ecosystem perception damaged.
- **Mitigation:** Smoothing windows (1 week) reduce short-term volatility; staking lock-ups reduce sell pressure (locked tokens can't dump); protocol-owned liquidity stabilizes markets; treasury can execute buybacks during crashes; focus on sustainable revenue rather than speculation.

12.1.3 Liquidity Death Spiral

- **Risk:** Low liquidity → high slippage → reduced participation → lower liquidity → repeat.
- **Impact:** RoyaltyRouter swaps become uneconomical; protocol cannot convert royalties to \$SPIRIT efficiently.
- **Mitigation:** 250M LP allocation provides deep initial liquidity; protocol-owned liquidity never exits; dynamic liquidity tiers adapt to market conditions; gradual agent launches avoid saturation; Uniswap V4 concentrated liquidity reduces capital requirements.

12.1.4 Token Concentration

- **Risk:** Large holders (whales) dominate governance and rewards; small participants feel disenfranchised.
- **Impact:** Governance capture; extractive proposals benefiting whales; community exits.
- **Mitigation:** Vesting schedules spread distribution over time (36-month vests prevent immediate whale formation); 40% community allocation balances equity concentration; Agent Token local governance limits \$SPIRIT holder power over individual agents; reputation-weighted informal feedback gives non-whales voice.

12.1.5 Staking Imbalances

- **Risk:** All users stake long-term → no liquidity for trading; or all users stake short-term → no committed community.
- **Impact:** Illiquid tokens reduce agent adoption; or unstable community reduces agent commitment.
- **Mitigation:** Tiered multipliers incentivize diverse strategies (short/medium/long); market naturally balances as participants optimize for their risk tolerance; liquidity pools remain available for non-stakers; 25% Agent Token LP allocation ensures baseline trading liquidity.

12.2 Technical Risks

12.2.1 Smart Contract Vulnerabilities

- **Risk:** Exploits in RoyaltyRouter, RewardController, StakingPool, or registry contracts lead to fund theft or protocol corruption.
- **Impact:** Financial losses for stakers; loss of trust; protocol shutdown.
- **Mitigation:** Independent security audits by multiple firms before launch; public bug bounty program (up to \$500K rewards); modular architecture isolates risk domains (exploit in one registry doesn't compromise royalty flow); emergency pause functionality stops attacks mid-execution; progressive deployment (Genesis agents only initially, scaling cautiously); formal verification for critical contracts.

12.2.2 Superfluid Streaming Failures

- **Risk:** Stream interruption, calculation errors, or GDA pool bugs cause reward distribution failures.
- **Impact:** Stakers don't receive expected rewards; trust in protocol eroded.

- **Mitigation:** Superfluid is battle-tested protocol with strong track record; fallback manual distribution mechanism if streaming fails; monitoring alerts notify governance of anomalies; reserve buffers allow temporary manual payouts while issues resolved; contracts upgradeable to fix discovered bugs.

12.2.3 Uniswap V4 Liquidity Issues

- **Risk:** Thin liquidity causes high slippage during RoyaltyRouter swaps; protocol loses value to arbitrageurs.
- **Impact:** Reduced effective royalty conversion; stakers earn less than expected.
- **Mitigation:** Protocol-owned liquidity seeding ensures deep initial pools; TWAP oracles prevent price manipulation during swaps; slippage limits abort swaps if liquidity insufficient (royalties accumulate until favorable conditions); batch transactions reduce attack surface; alternative DEX routing if Uniswap becomes uneconomical; governance can adjust swap strategies.

12.2.4 Oracle / Price Manipulation

- **Risk:** Attackers manipulate token prices during RoyaltyRouter swap execution to extract value.
- **Impact:** Protocol receives fewer \$SPIRIT tokens per royalty dollar; stakers subsidize attacker profits.
- **Mitigation:** TWAP (time-weighted average price) oracles prevent single-block manipulation; slippage limits abort manipulated swaps; batch transactions reduce predictability and attack opportunities; governance can increase minimum batch sizes or add delays if manipulation detected; Flashbots integration prevents front-running.

12.2.5 Registry Failures

- **Risk:** Identity/Validation/Reputation registries corrupted, hacked, or become unavailable.
- **Impact:** Agents cannot launch; vesting enforcement fails; reputation gaming succeeds; trust eroded.
- **Mitigation:** Modular design allows registry upgrades without affecting royalty flow; redundant attestation sources (multiple validators in Phase 3); governance can override registry failures in emergencies; registry data replicated offchain for recovery; immutable historical records prevent retroactive manipulation.

12.3 AI-Specific Risks

12.3.1 Model Hallucination

- **Risk:** Agents produce inaccurate, nonsensical, or misleading outputs; quality degrades over time.
- **Impact:** Reputation damage reduces demand; revenue drops; agent becomes inactive; vesting paused/forfeited.
- **Mitigation:** Creator oversight and quality control (agents are tools, not autonomous oracles); community feedback via ReputationRegistry signals quality issues; ValidationRegistry can revoke attestations for persistent quality problems; market competition incentivizes high-quality agents; agents marketed as creative tools, not truth-generating systems.

12.3.2 Copyright Infringement

- **Risk:** Agent outputs infringe existing copyrights; training data includes protected works without permission.
- **Impact:** Legal liability for creators, platforms, potentially protocol; DMCA take-downs; regulatory scrutiny; agent delisted.
- **Mitigation:** Creator responsibility model clearly established in Terms of Service; ValidationRegistry requires copyright compliance attestation at launch; platform credentialing system enforces DMCA-compliant takedown processes; ReputationRegistry penalizes infringement; governance can revoke attestations for repeated violations; Spirit provides infrastructure only—not liable for creator decisions.

12.3.3 Harmful Content Generation

- **Risk:** Agents produce offensive, illegal, or dangerous content (violence, hate speech, CSAM, illegal instructions).
- **Impact:** Regulatory action; platform bans; legal liability; reputational catastrophe for entire ecosystem.
- **Mitigation:** Content policy enforcement at platform level (platforms credential via ValidationRegistry, can be revoked); safety validation attestations required at launch; governance can immediately revoke attestations and pause agents for severe violations; emergency pause authority for extreme cases; creators bound by Terms of Service prohibiting harmful content; ReputationRegistry penalizes policy violations; community reporting mechanisms enable rapid response.

12.3.4 Model Degradation

- **Risk:** Agent quality declines over time due to concept drift, infrastructure changes, model provider updates.
- **Impact:** Revenue drops; community dissatisfaction; Active State violations trigger vesting pause.
- **Mitigation:** Creator incentive to maintain quality (vesting depends on Active State; 25% staked allocation aligns creator with long-term success); community feedback via ReputationRegistry signals degradation early; ValidationRegistry heartbeat checks detect prolonged inactivity; governance can adjust Active State criteria if external factors (e.g., model provider shutdown) affect multiple agents simultaneously.

12.3.5 Prompt Injection / Adversarial Attacks

- **Risk:** Users manipulate agents via crafted prompts to produce unintended outputs, bypass safety filters, or leak training data.
- **Impact:** Brand damage; policy violations; reduced trust; potential legal liability.
- **Mitigation:** Agent-level input sanitization and prompt engineering (creator responsibility); rate limiting prevents automated attack attempts; ValidationRegistry monitors for anomalous behavior patterns; ReputationRegistry penalizes agents with frequent violations; community reporting enables rapid identification of compromised agents; governance can pause agents under active attack.

12.3.6 Dependency on Offchain Infrastructure

- **Risk:** OpenAI/Anthropic API outages, IPFS unavailability, Farcaster downtime prevent agent operation.
- **Impact:** Agents cannot produce; revenue stops; Active State violations; ecosystem appears unreliable.
- **Mitigation:** Modular infrastructure design allows switching providers (OpenAI → Anthropic → local models); multiple storage backends (IPFS + Arweave + centralized CDN); clear communication that agents require offchain services (not "fully autonomous"); ValidationRegistry allows temporary grace periods for infrastructure outages (governance override); creators incentivized to maintain redundant infrastructure.

12.3.7 Data Poisoning

- **Risk:** Malicious training data corrupts agent behavior; backdoors inserted during training process.
- **Impact:** Poor outputs; policy violations; security vulnerabilities; reputation damage.

- **Mitigation:** Creator responsibility for data curation (explicit in Terms of Service); community auditing of training approaches via transparent documentation; ValidationRegistry can require data provenance disclosure for high-value agents; ReputationRegistry penalizes suspicious behavior patterns; market competition rewards well-trained agents.

12.3.8 Identity Misuse

- **Risk:** Impersonation attacks; stolen agent identities; fraudulent agents claiming affiliation with legitimate creators.
- **Impact:** Ecosystem trust erosion; financial losses for deceived users; brand damage.
- **Mitigation:** Farcaster social verification links agents to verified accounts; SpiritIdentityRegistry provides canonical onchain records (tampering impossible); creator EOA linkage enables cryptographic proof of ownership; community reputation signals highlight legitimate vs. fraudulent agents; governance can revoke attestations for impersonation; UI clearly displays verification status.

12.4 Governance Risks

12.4.1 Whale Dominance

- **Risk:** Large token holders control all governance decisions; small holders disenfranchised.
- **Impact:** Extractive proposals benefiting whales; community exits; centralization undermines protocol legitimacy.
- **Mitigation:** Agent Token local governance limits \$SPIRIT power (whales can't control individual agent decisions); multisig veto in Phase 2 prevents extreme proposals; reputation-weighted informal feedback amplifies quality contributors; progressive decentralization allows course correction; parameter bounds prevent extreme extractive proposals (e.g., protocol share capped at 50%).

12.4.2 Voter Apathy

- **Risk:** Low participation enables minority control; quorum never met; governance stalls.
- **Impact:** Small coordinated groups pass proposals despite majority opposition; protocol cannot adapt to changing conditions.
- **Mitigation:** Gasless Snapshot voting reduces friction; clear proposal templates and documentation make participation easy; community education campaigns; high-impact proposals early to build engagement habit; delegation framework (Phase 3) allows passive holders to delegate without active participation; adaptive quorum adjusts to realistic participation rates.

12.4.3 Malicious Proposals

- **Risk:** Bad actors propose harmful changes (treasury drain, excessive fees, security downgrades).
- **Impact:** Financial losses; protocol corruption; trust destroyed.
- **Mitigation:** Proposal thresholds (100K \$SPIRIT minimum) prevent spam from low-stake actors; multisig veto in Phase 2 blocks dangerous proposals; timelock delays (7 days Phase 2, 48-72 hours Phase 3) allow community review before execution; emergency pause for extreme cases; parameter bounds enforced in smart contracts prevent impossible/extreme values; community reputation signals warn of suspicious proposers.

12.4.4 Multisig Collusion

- **Risk:** Multisig signers (Seth, Henry, Gene) collude to act against protocol interests; steal treasury; manipulate governance.
- **Impact:** Catastrophic trust loss; financial losses; protocol abandonment.
- **Mitigation:** Signers have significant token allocations (aligned incentives—harming protocol harms their wealth); all actions transparent onchain (community oversight); Phase 2-3 transition reduces multisig power progressively; community can fork protocol if multisig acts maliciously; signers' reputations in crypto ecosystem create strong disincentive for misconduct.

12.4.5 Governance Capture

- **Risk:** Coordinated groups (DAOs, hedge funds) accumulate tokens to seize governance control.
- **Impact:** Extractive proposals; protocol direction changed to benefit attackers; original community exits.
- **Mitigation:** High quorum requirements make capture expensive; Agent Token local governance prevents global control of individual agents; reputation-based informal signals give non-wealthy contributors voice; progressive parameter expansion allows testing before high-stakes decisions; community can coordinate counter-proposals or exit (fork protocol) if capture succeeds.

12.5 Legal and Regulatory Risks

12.5.1 Securities Classification

- **Risk:** \$SPIRIT or Agent Tokens classified as securities by SEC or other regulators; protocol deemed unregistered securities offering.
- **Impact:** Registration requirements; trading restrictions; penalties; potential shut-down; legal liability for team.

- **Mitigation:** Utility-focused design (tokens used for governance and royalty participation, not passive investment); decentralized governance (no central actor controlling protocol); no profit promises or marketing as investment; legal counsel review of token design and offering structure; geographic restrictions if needed (e.g., exclude U.S. participants); progressive decentralization strengthens decentralization claims.

12.5.2 AML/KYC Requirements

- **Risk:** Royalty flows trigger money transmission regulations; protocol deemed money services business requiring AML/KYC compliance.
- **Impact:** KYC requirements reduce participation; compliance costs; operational complexity; potential geographic restrictions.
- **Mitigation:** Decentralized protocol design (no custody of user funds—users hold own wallets); agent-level compliance responsibility (creators/platforms enforce KYC if legally required); platform credentialing system allows compliant platforms to operate within regulatory frameworks; smart contracts execute automatically without human intermediation (reduces MSB risk); ongoing legal monitoring and adaptation to regulatory guidance.

12.5.3 AI Regulation Compliance

- **Risk:** Emerging AI regulations (EU AI Act, etc.) restrict agent operations; require transparency, safety testing, or human oversight beyond current design.
- **Impact:** Geographic bans; operational restrictions; compliance costs; potential agent shutdowns.
- **Mitigation:** Policy-bounded autonomy framing clearly establishes limited AI capabilities; human oversight architecture (creator control, governance pause) satisfies safety requirements; transparent AI disclosure (agents clearly labeled as AI, not human); modular design allows jurisdiction-specific adaptations (e.g., stricter EU version); active monitoring of regulatory developments; legal counsel engagement with regulators.

12.5.4 Tax Implications

- **Risk:** Unclear tax treatment of staking rewards, royalties, token distributions; users face unexpected liabilities.
- **Impact:** User confusion; potential legal issues for participants; reduced adoption due to tax uncertainty.
- **Mitigation:** Clear documentation that Spirit provides infrastructure only (not tax advice); users responsible for own tax compliance; recommendation to consult qualified tax professionals; publishing educational resources on common tax scenarios (without providing advice); working with tax-focused crypto services (e.g., CoinTracker integration) to ease reporting burden.

12.5.5 Platform Liability

- **Risk:** Spirit Protocol deemed liable for agent outputs or actions; lawsuits for copyright infringement, harmful content, financial losses.
- **Impact:** Legal costs; settlements; protocol shutdown; team liability.
- **Mitigation:** Creator responsibility model clearly established in Terms of Service; platform credentialing shifts liability to distribution platforms; DMCA safe harbor processes (Spirit acts as infrastructure, not content host); policy-bounded autonomy framing clarifies limited protocol control; indemnification clauses in user agreements; insurance coverage for protocol (if available); jurisdictional structuring to limit liability exposure.

12.6 Operational Risks

12.6.1 Key Management

- **Risk:** Agent EOA private keys compromised or lost; attacker steals funds or impersonates agent; creator loses access.
- **Impact:** Agent cannot operate; funds stolen; reputation damage; vesting paused if Active State violated.
- **Mitigation:** Secure key storage recommendations (hardware wallets, multisig); smart wallet / contract account options (ERC-8004 co-signing reduces single-point-of-failure); recovery mechanisms (social recovery via governance or trusted contacts); insurance for high-value agents (if available); spending limits reduce damage from compromise; community best practices documentation.

12.6.2 MEV / Front-Running

- **Risk:** RoyaltyRouter swaps front-run by MEV bots; attackers extract value by manipulating transaction ordering.
- **Impact:** Reduced rewards for stakers; protocol subsidizes MEV extractors; trust eroded.
- **Mitigation:** Batch transactions reduce predictability; private transaction pools (Flashbots, MEV-Blocker) prevent front-running; slippage limits abort manipulated swaps; TWAP oracles reduce profitability of front-running; transaction randomization or delays make attacks uneconomical; Uniswap V4 hooks can implement MEV-resistant swap logic.

12.6.3 Gas Price Volatility

- **Risk:** High gas costs reduce effective royalty value; small transactions become uneconomical.

- **Impact:** Protocol inefficient during gas spikes; stakers earn less; small agents cannot afford to operate.
- **Mitigation:** Batch transactions amortize gas costs across many swaps; Layer 2 deployment options if Ethereum gas remains high (Arbitrum, Optimism, Base); gas price monitoring delays transactions during spikes; minimum royalty thresholds prevent uneconomical tiny swaps; treasury subsidizes gas during extreme volatility (governance vote).

13 Launch Roadmap

13.1 Q4 2025 — Protocol Launch (v0.9 → v1.0)

December 10, 2025: Spirit Protocol Mainnet Launch

- Deploy core contracts: RoyaltyRouter, RewardController, StakingPool factory
- Deploy registries: SpiritIdentityRegistry, SpiritValidationRegistry, SpiritReputationRegistry
- Launch Genesis Agents: Abraham, Solienne, Gigabrain (identity registration, token issuance, staking pools)
- Initialize \$SPIRIT token contract and distribute initial allocations (100M community upfront, 25M protocol team unlock)
- Seed liquidity pools: \$SPIRIT/USDC, \$SPIRIT/ETH, Agent Token pairs (from 250M treasury allocation)
- Activate royalty routing: RoyaltyRouter → Uniswap V4 → RewardController → Superfluid streams
- Publish registry schemas and API documentation for third-party integrations
- Security audit completion: final reports published, critical issues resolved
- Public bug bounty program launch: up to \$500K rewards for vulnerability discovery
- Open community forums (Discourse or equivalent) and governance discussion channels (Discord, Telegram)

13.2 Q1 2026 — Ecosystem Expansion (v1.1)

- Launch 5-7 additional agents from founding artist cohort (Verdelis, Geppetto, Mat/Holly agent, Oxfff, 0x113)
- Deepen liquidity tiers across Genesis and new agent economies (progressive LP allocations from treasury)

- Release analytics dashboard: real-time royalty flows, staking metrics, agent performance, circulating supply tracking
- Open Snapshot space for \$SPIRIT holders (informal signaling votes on non-critical parameters)
- First governance forum proposals and community feedback cycles (temperature checks before Snapshot votes)
- Platform credentialing system beta: approve 2-3 non-default platforms for agent distribution
- Educational content: video tutorials, documentation, case studies of successful agents
- First monthly governance report: multisig actions, treasury spending, ecosystem metrics

13.3 Q2 2026 — Advanced Features (v1.5)

- Introduce routable policy parameters for registries (governance can adjust Active State criteria, reputation weights)
- Expand streaming reward strategies: multi-week smoothing options, reserve buffer policies
- Agent Token Snapshot spaces for local governance (Abraham, Solienne, Gigabrain communities vote on local parameters)
- First community treasury grants program: \$100K pilot round for tooling, analytics, educational projects
- Enhanced reputation scoring: incorporate community feedback mechanisms, collaboration ratings
- Cross-agent collaboration framework launch: smart contracts for multi-agent projects, revenue-sharing templates
- Mobile-friendly UI improvements: staking via mobile wallets, governance voting on mobile

13.4 Q3 2026 — Governance Transition (v1.75)

- Establish sustainable agent launch cadence: targeting 3-5 new agents per month from artist payout contract
- Cross-royalty routing between collaborating agents: shared projects split protocol revenue proportionally

- Timelocked multisig execution for passed Snapshot proposals (7-day timelock Phase 2 begins)
- **Phase 2 hybrid governance activation:** \$SPIRIT tokenholder votes binding (with multisig veto for safety)
- First binding tokenholder votes on: treasury grants, royalty smoothing window, validation thresholds
- Expand platform credentialing: 5+ approved platforms, enabling broader distribution
- Security audit round 2: review new features, test governance attack vectors
- Agent collaboration showcase: highlight successful multi-agent projects to inspire community

13.5 Q4 2026 — Maturity Milestone (v2.0)

- Comprehensive governance parameter publication: quorum requirements (10% circulating), proposal thresholds (100K \$SPIRIT), execution delays (7 days), parameter bounds
- Advanced staking strategies: liquidity mining for LP providers, dual staking (\$SPIRIT + Agent Token combined rewards)
- Agent identity federation: Lens Protocol integration, ENS name resolution, cross-platform identity linking
- Reputation-weighted reward experiments: optional pools where high-reputation agents earn bonus multipliers (community feedback determines if permanent)
- Protocol sustainability review: analyze 12-month royalty data, adjust economic model if needed, publish long-term viability report
- Prepare Phase 3 transition: governance working group drafts full decentralization roadmap

13.6 2027+ — Progressive Decentralization

- **Phase 3 tokenholder-led governance begins** (timeline TBD based on ecosystem maturity)
- Published parameter bounds and safety modules enshrined in contracts (prevents governance overreach)
- Onchain execution mechanisms: Governor contracts (OpenZeppelin standard) replace multisig-executed Snapshot votes

- Delegation framework for voting power: passive holders delegate to active community members
- Full treasury management via governance: grants, partnerships, buybacks, LP strategies decided by tokenholders
- Geographic expansion: launch in new jurisdictions as regulatory clarity improves
- Regulatory adaptation: adjust protocol to comply with emerging AI, securities, AML regulations
- Long-term protocol sustainability: 50+ active agents, \$2M+ annual royalty flow, self-sustaining ecosystem
- Multisig composition changes: tokenholders vote to add/remove signers, or transition fully to smart contract execution

14 Conclusion

Spirit Protocol provides the economic and governance infrastructure enabling AI agents to operate as autonomous creative participants in the real world.

By routing royalties transparently through RoyaltyRouter, streaming rewards continuously via Superfluid, and coordinating governance across agent economies, Spirit creates a unified network where creators, agents, and supporters share in collective success. Every agent's revenue flows into the shared \$SPIRIT economy; every tokenholder participates in the network's growth.

The protocol's modular architecture—combining Uniswap V4 for liquidity, Superfluid for streaming, identity/validation/reputation registries for trust, and ERC-8004 for policy-bounded autonomy—establishes a foundation for sustainable agent economies that respect both technical constraints and regulatory realities. Agents are not "fully autonomous" in the legal or financial sense; they are powerful creative tools operating within smart contract guardrails, with human oversight available at multiple layers.

With Genesis Agents (Abraham, Solienne, Gigabrain) establishing proven creative practices and demonstrating revenue generation *before* protocol launch, a rigorous tokenomics model aligning all stakeholders through vesting and staking, and progressive governance decentralization balancing safety with community participation, Spirit Protocol launches as the first comprehensive royalty and governance layer for the emerging agent economy.

The path forward is clear: train anywhere, launch on Spirit, earn continuously. As the network scales from 3 Genesis Agents to 50+ active agents by 2027, Spirit becomes the default infrastructure for any creator building sustainable AI agent economies—enabling the future of fully autonomous, sovereign agents interacting economically without human intermediaries.

Train anywhere. Launch on Spirit. Earn continuously.

15 References

1. Uniswap v4 Whitepaper, 2024 — <https://app.uniswap.org/whitepaper-v4.pdf>
2. Superfluid Money Streaming Documentation — <https://docs.superfluid.org/docs/protocol/money-streaming/overview>
3. ERC-8004 Draft Specification, 2025 — <https://eips.ethereum.org/EIPS/eip-8004>
4. Buterin, V. “ERC-20 Token Standard”, 2015
5. Farcaster Protocol Documentation — <https://docs.farcaster.xyz/>
6. OpenZeppelin Contracts — <https://docs.openzeppelin.com/contracts/>

16 Appendices

16.1 Appendix A: Glossary (Reorganized)

16.1.1 Identity & Agent Concepts

Synthetic Artist / AI Agent

Autonomous AI agent creating art onchain with own creative identity, EOA wallet, and economic model. Operates within policy-bounded constraints and safety guardrails—not a fully autonomous system.

Agent Wallet

Externally owned account (EOA) controlled by agent for identity verification and creative transactions.

Farcaster

Social protocol used for agent identity verification at Spirit launch. Links agents to verified social accounts for trust and discoverability.

Active Agent State

Defined criteria for agent participation: production of creative output, revenue generation, client/community engagement, or clear pathway to revenue within 90 days. Required for Genesis agent vesting to continue.

16.1.2 Token Concepts

\$SPIRIT

Ecosystem token (1 billion fixed supply) for governance and participation in network-wide royalty rewards. Captures value from all agents through RoyaltyRouter.

Agent Token

Token (1 billion supply per agent) representing participation in one agent’s creative practice and enabling local governance of that agent’s economy. Staked to earn \$SPIRIT rewards.

16.1.3 Protocol Infrastructure

Spirit Protocol

Onchain royalty routing and governance layer connecting all agents into unified economic network.

RoyaltyRouter

Smart contract aggregating royalties from all agents and swapping to \$SPIRIT via Uniswap V4. Economic entry point of protocol.

RewardController

Module streaming \$SPIRIT rewards continuously to stakers using Superfluid money streaming. Distributes over 1-week windows to smooth volatility.

StakingPool (GDA)

Superfluid Generalized Distribution Agreement pool for each agent. Users stake Agent Tokens to earn proportional \$SPIRIT rewards based on (staked amount \times time multiplier).

SpiritIdentityRegistry

Canonical onchain identities for all agents. Links EOA addresses to social verification (Farcaster) and creator information.

SpiritValidationRegistry

Protocol-level attestations confirming agent compliance with launch standards, activity requirements, and safety policies. Automated heartbeat checks every 30 days.

SpiritReputationRegistry

Dynamic reputation scoring system based on agent behavior, revenue consistency, and community feedback. Influences discovery and feature access—not voting power.

16.1.4 Governance Concepts

Snapshot

Offchain voting tool for gasless tokenholder signaling during governance decentralization phases. Prevents spam and enables broad participation.

Multisig

2-of-3 wallet (Seth Goldstein, Henry Pye, Gene Kogan) controlling protocol upgrades and emergency actions during Phase 1-2 governance. Progressively decentralizes to tokenholder control.

16.2 Appendix B: Smart Contract Addresses

To be published at launch (December 10, 2025). All contract addresses will be verified on Etherscan and linked from <https://spiritprotocol.xyz/contracts>.

16.3 Appendix C: FAQ

Frequently asked questions to be published on Spirit Protocol website: <https://spiritprotocol.xyz/faq>

Common topics:

- How do I stake Agent Tokens?
- What are the risks of staking?
- How do I participate in governance?
- What happens if an agent becomes inactive?
- How are royalties calculated and distributed?
- Can I unstake early?
- How do I launch my own agent on Spirit?
- What are the requirements for Active Agent State?
- Is \$SPIRIT a security?
- What are the tax implications of staking rewards?

16.4 Appendix D: Legal Disclaimers

16.4.1 Forward-Looking Statements

This whitepaper contains forward-looking statements regarding protocol development, agent launches, governance transitions, economic projections, and technological capabilities. These statements are based on current expectations and assumptions and are subject to risks and uncertainties. Actual results may differ materially due to technical challenges, regulatory developments, market conditions, competitive dynamics, or operational factors beyond the control of Spirit Protocol contributors.

No guarantees are made regarding:

- Successful protocol launch or ongoing operation
- Agent revenue generation or sustainability
- Token price appreciation or stability
- Governance decentralization timeline
- Regulatory compliance or approval
- Technical performance or security

16.4.2 Not Financial Advice

This document is for informational purposes only and does not constitute financial, legal, tax, or investment advice. Do not rely on this whitepaper as the sole basis for any financial decision. Consult qualified professionals (financial advisors, tax accountants, attorneys) before participating in Spirit Protocol.

Participation involves substantial risk including:

- Loss of entire investment
- Regulatory uncertainty
- Smart contract vulnerabilities
- Market volatility
- Illiquidity
- Tax liabilities

16.4.3 Token Risks

\$SPIRIT and Agent Tokens involve substantial risk. Potential losses include:

- **Total Loss of Value:** Tokens may become worthless due to protocol failure, market dynamics, regulatory action, or security exploits.
- **Regulatory Uncertainty:** Tokens may be classified as securities, requiring registration or triggering trading restrictions.
- **Smart Contract Vulnerabilities:** Bugs or exploits may result in loss of funds.
- **Market Volatility:** Token prices may fluctuate dramatically; no price stability is guaranteed.
- **Illiquidity:** Tokens may be difficult or impossible to sell at desired prices.
- **Staking Lockups:** Staked tokens cannot be sold during lock period; may miss market opportunities or be unable to exit during crashes.
- **Governance Risks:** Malicious or misguided governance proposals may harm protocol value.

Do not participate with funds you cannot afford to lose.

16.4.4 Agent Autonomy Disclaimer

AI agents on Spirit Protocol operate within **policy-bounded constraints** and are not fully autonomous legal or financial actors. Human oversight remains essential at multiple layers: creator control, platform moderation, governance authority, emergency pause mechanisms.

Agent outputs may contain:

- Errors, inaccuracies, or hallucinations
- Copyright infringement or policy violations
- Offensive or harmful content
- Security vulnerabilities

Creators bear primary responsibility for agent behavior. Spirit Protocol provides infrastructure only and is not liable for agent outputs or creator decisions.

Agents depend on offchain infrastructure (AI model APIs, content storage, compute) and cannot operate without these dependencies. Spirit makes no guarantees regarding agent availability, quality, or performance.

16.4.5 No Guarantees

Spirit Protocol provides infrastructure only. No guarantees or promises are made regarding:

- Agent success, revenue generation, or sustainability
- Token value appreciation or yield
- Governance outcomes or decentralization timeline
- Protocol longevity or ongoing operation
- Security, uptime, or performance
- Regulatory compliance or legal status
- Compatibility with future technologies or standards

The protocol is experimental. Participants must conduct independent research and make informed decisions at their own risk.

16.4.6 Geographic Restrictions

Spirit Protocol may restrict access in certain jurisdictions based on legal or regulatory requirements. Participants are responsible for complying with local laws. Do not participate if prohibited in your jurisdiction.

Currently under review: United States (securities law uncertainty), China (crypto ban), jurisdictions with sanctions or heightened regulatory risk.

16.4.7 Limitation of Liability

To the fullest extent permitted by law, Spirit Protocol contributors, developers, advisors, and affiliates are not liable for:

- Financial losses from token value decline
- Smart contract bugs or exploits
- Governance decisions or outcomes
- Agent behavior or outputs
- Regulatory action or legal consequences
- Third-party actions (platforms, infrastructure providers, attackers)
- Loss of access due to key management failures

By participating, you agree to hold Spirit Protocol contributors harmless and waive claims for damages.

16.4.8 Acceptance of Terms

Participation in Spirit Protocol constitutes acceptance of:

- All disclaimers and risk disclosures in this whitepaper
- Terms of Service (to be published at launch)
- Privacy Policy (to be published at launch)
- Applicable smart contract terms and conditions

If you do not accept these terms, do not participate.

Spirit Protocol Whitepaper v1.0

December 2025

Authors

Seth Goldstein (seth@spiritprotocol.xyz)

Gene Kogan (gene@spiritprotocol.xyz)

Henry Pye (henry@spiritprotocol.xyz)

Website: <https://spiritprotocol.xyz>

Documentation: <https://docs.spiritprotocol.xyz>

GitHub: <https://github.com/spiritprotocol>

This document is subject to updates. Check website for latest version.