



# UMBRAE

---

## The Institutional Bridge for Decentralized Finance

---

Strategic Business Plan & Ecosystem Outlook v4.0

Ignis AI Labs | March 2026

---

### 1. Executive Summary

---

Umbrae is a Multi-Chain Trading Infrastructure Layer developed by Ignis AI Labs. We are solving the fragmentation of DeFi by building a unified “DeFi-CeFi Hybrid” ecosystem. Our platform bridges the gap between the transparency of decentralized finance and the capability of centralized institutions.

**As of Q1 2026, Umbrae is live.** The platform is operational on Base (Ethereum L2), our SAGE AI engine is serving users, and our fee infrastructure is generating revenue. This document supersedes all prior versions of the business plan and reflects Umbrae's transition from pre-launch positioning to an operational platform with a clear, milestone-driven expansion path.

While our core trading rails remain permissionless, we overlay a proprietary compute layer to offer features — such as AI-driven automation and institutional liquidity management — that pure DeFi cannot support. This hybrid model allows us to service everyone from the retail novice to the compliance-focused whale.

#### 1.1 Strategic Differentiators

- ▲ **The First EVM-Native DLMM:** Umbrae built and deployed the first Dynamic Liquidity Market Maker of its kind on an EVM chain. Not a fork of existing infrastructure — engineered from the ground up, audited, and penetration tested. Because it is EVM-native, the architecture is portable to any EVM-compatible blockchain — Base today, Arbitrum, Optimism, BNB Chain, Avalanche, Hyperliquid, and any future EVM chain tomorrow. This is not a single-chain product; it is a deployable standard.
- ▲ **Universal Infrastructure:** A single, unified interface across EVM chains, with Base as our launch chain. The same contracts, the same routing, the same fee structure — deployed wherever EVM runs. Solana support will follow the EVM rollout.
- ▲ **The Hybrid Model:** A bifurcated access system. Layer 1 is open, permissionless DeFi (Swaps, LP). Layer 2 is a gated, KYC-compliant environment for advanced features like AI-managed strategies and automated rebalancing.
- ▲ **SAGE AI:** The Strategic Alpha Generation Engine — now live — offering institutional-grade analysis via a proprietary credit-based compute system. SAGE currently operates as an intelligent routing layer, channeling best-in-class external models through a unified interface optimized for on-chain analysis. A proprietary model trained on Umbrae platform data is on the roadmap.
- ▲ **Sovereign Trading Infrastructure:** A proprietary arbitrage system that generates real volume, keeps pool prices accurate with market, and distributes profits to NFT holders. Most protocols depend on external arbitrageurs — we built our own.
- ▲ **Revenue Alignment:** A contractual commitment to distribute 40% of all platform fees back to the community (NFT and Token holders), enforced at the smart contract level.

## 1.2 At-A-Glance: Key Constraints & Metrics

Metric	Specification	Constraint
Max Token Supply	5,000,000 \$U1	Fixed (Non-inflationary)
Max NFT Licenses	777 Ignis Elite	Fixed (Never increasing)
Community Fee Share	40% of Platform Fees	Perpetual (Smart Contract enforcement)
Live Chain	Base (EVM)	Operational Q1 2026
SAGE AI	v1 Live	Credit system operational

## 2. The Market Opportunity

### 2.1 The Problem: Fragmentation & Friction

The current landscape is hostile to serious capital:

1. **Fragmentation:** Traders must navigate disjointed tools for different chains, fracturing liquidity and attention.
2. **Safety:** The space is flooded with scams, rug pulls, and low-quality listings with no verification standards.
3. **Inefficiency:** Retail traders face high slippage and MEV attacks while competing against institutional bots with no defense.

### 2.2 The Umbrae Solution: Unified Architecture

We provide a comprehensive ecosystem that services both the Trader and the Project Issuer.

- ▲ **For Traders:** A gamified, RPG-like trading experience backed by enterprise-grade routing and AI assistance.
- ▲ **For Projects:** A "Market Making as a Service" infrastructure. We allow legitimate projects to launch across multiple chains with professional liquidity structures from Day 1.

## 3. Platform Vision: The Hybrid Thesis

Umbrae is designed as a DeFi-CeFi Hybrid. This is a deliberate architectural choice to ensure long-term viability and regulatory resilience.

### 3.1 The Two-Tiered Access Model

Certain high-performance features cannot be made permissionless due to computational limits and regulatory requirements regarding fund management.

#### Layer 1: The Open Protocol (Permissionless)

- ▲ Access: Open to all (No KYC).
- ▲ Features: Multi-Chain Swaps, Manual DLMM Liquidity Provisioning, Portfolio Tracking, MEV Protection.
- ▲ Philosophy: Self-custody is absolute.

## Layer 2: The Institutional Layer (Gated)

- ▲ Access: Requires KYC / Identity Verification (Ignis Elite).
- ▲ Features: Automated DLMM Rebalancing, AI-Managed Strategies, Copy-Trading, and supported bot deployment on Umbrae infrastructure (arbitrage, market making, custom strategies).
- ▲ Rationale: The gate serves two purposes. First, features involving active management of strategy or high-frequency compute resources require regulatory compliance to ensure platform longevity. Second — and more fundamentally — the KYC layer is a sovereignty filter. We must know who we are working with to ensure that no actor enters the ecosystem who would compromise decentralized sovereignty. The level of intelligence and vetting we apply to this process is designed to rival that of any institution or state-level apparatus, because protecting the integrity of this network demands nothing less.

---

## 4. Technology & Capabilities

### 4.1 SAGE: Strategic Alpha Generation Engine — Live

SAGE is our flagship AI interface, now operational and serving users on the platform.

- ▲ **Current Architecture:** SAGE operates as an intelligent routing layer — channeling best-in-class external AI models through a unified interface purpose-built for on-chain analysis. This approach delivers immediate, high-quality output while keeping infrastructure lean. Users interact with a single interface; SAGE handles model selection, context injection, and on-chain data integration behind the scenes.
- ▲ **The Vision:** Train a proprietary model on Umbrae platform data — trading patterns, liquidity dynamics, pool performance, and verification outcomes. This is a milestone on the roadmap, gated by sufficient platform-native data to make the model meaningfully better than general-purpose alternatives.
- ▲ **The Credit System:** SAGE operates on a “Compute Credit” economy, now live.
  - ▲ *Basic:* Simple price queries and trend checks (Free/Low cost).
  - ▲ *Advanced:* Complex requests (e.g., “Backtest a grid-trading strategy on SOL/USDC for the last 3 months”) consume credits.
  - ▲ *Revenue:* Credit purchases contribute to the platform fee pool, driving yield for NFT/Token holders.

### 4.2 Infrastructure: Current State

We utilize enterprise-grade hardware to ensure low latency and high uptime.

- ▲ **Compute:** Dual AMD EPYC 9754 (256 cores) for parallel trade processing and SAGE compute.
- ▲ **GPU:** NVIDIA RTX 6000 for AI inference and compute-intensive operations.
- ▲ **Memory:** 512GB DDR5-4800MHz ECC to handle peak load state management.
- ▲ **Speed:** 80–120ms response time (vs. industry standard 1000ms+).
- ▲ **RPC Nodes:** Private EVM RPC node infrastructure is actively being set up.
- ▲ **Security:** Multi-signature treasury; active audit relationships with top-tier firms.
  - ▲ *Completed:* Hashlock — Penetration test on Telegram/Discord wallet infrastructure (August 2025). Rating: Secure.
  - ▲ *Completed:* Hashlock — Follow-up audit for expanded functionality.

### 4.3 Infrastructure: The Build-Out Vision

As revenue scales, we are committed to building fully private, owned infrastructure. This is not a cost-cutting measure — it is a competitive moat. **Private RPC node deployment is currently underway as an active priority.**

#### Why Private Infrastructure Matters:

- ▲ **Performance Edge:** Cloud providers (AWS, GCP) introduce latency and share resources. Private hardware means dedicated throughput with no noisy neighbors.
- ▲ **RPC Independence:** Running our own EVM RPC nodes eliminates dependency on third-party providers and removes rate limits that throttle competitors.
- ▲ **AI Sovereignty:** Self-hosted SAGE instances mean no API costs, no external rate limits, and full control over model fine-tuning.
- ▲ **Cost Efficiency:** At scale, owned hardware is 3–5x cheaper than cloud compute for sustained workloads.

#### The Target Stack:

Component	Purpose	Advantage
Private Server Racks	Core compute infrastructure	Dedicated resources, no shared tenancy
EVM RPC Nodes	Base/Arbitrum/Optimism access	Multi-chain independence
AI Compute Cluster	SAGE model hosting	Unlimited inference, custom fine-tuning
Redundant Networking	Multi-ISP failover	99.99% uptime guarantee

**Funding Source:** A portion of NFT sale proceeds and ongoing platform fees (via the Operations and Treasury allocations) is earmarked for infrastructure capital expenditure. Our leadership has direct experience in hardware architecture and data center buildouts — this is not outsourced.

### 4.4 EVM-First Multi-Chain Strategy

Umbrae's DLMM is the first of its kind built natively for EVM. This is a deliberate architectural decision — it means the entire infrastructure is portable to any EVM-compatible chain without rewriting contracts or re-auditing core logic.

- ▲ **Live:** Base (Ethereum L2) — operational as of Q1 2026.
- ▲ **Next:** Expansion across EVM chains. Arbitrum, Optimism, Ethereum mainnet, BNB Chain, Avalanche, and Hyperliquid are all on the expansion path. Each deployment leverages the same audited contracts with chain-specific configuration only.
- ▲ **Future:** Solana support will follow once EVM expansion is established. Base volume comes first.

This represents a strategic shift from the original plan. By building EVM-native from the start, we created a deployable standard rather than a single-chain product. Every new EVM chain that launches is a potential expansion target with minimal marginal cost. We are not locked to one ecosystem — we go where the volume is.

### 4.5 Proprietary Arbitrage Infrastructure — Live

Most DeFi protocols depend on external arbitrageurs and hope for organic volume. Umbrae built its own.

Our proprietary arbitrage bot is live on Base mainnet, running with real capital. It was deployed as a proof of concept, validated on-chain, and is now scaling. Every transaction is publicly verifiable.

#### How It Works:

The bot runs 24/7 scanning for price differences across DEXes on Base. It polls 5 DEXes simultaneously (Uniswap V3, PancakeSwap V3, Aerodrome, SushiSwap, and Umbrae DLMM), evaluating 20 routes per cycle every 3 seconds. When it identifies a price discrepancy, it executes a round-trip trade — buying where cheap, selling where expensive — in back-to-back transactions.

**Technical Capabilities:**

- ▲ Round-trip quote model — quotes exact buy and sell before committing capital.
- ▲ Dynamic trade sizing — automatically sizes positions based on available capital.
- ▲ Auto-recovery — if a sell leg fails, automatically reverses the buy to recover capital.
- ▲ Auto-rebalancing — detects portfolio skew and corrects without manual intervention.
- ▲ Cached RPC architecture — optimized to minimize infrastructure costs.
- ▲ Adaptive polling — fast (3s) when opportunities exist, slow (15s) when flat.

Fully automated. No frontend dependency. No manual oversight. Every layer is proprietary.

**Why This Matters for the Ecosystem:**

- ▲ Umbrae pools stay accurately priced with the broader market — because we enforce it.
- ▲ LPs earn real fees from real volume — because the bot generates it.
- ▲ The protocol earns revenue — because every trade carries a spread.
- ▲ NFT holders benefit directly — 35% of arbitrage bot profits are distributed to locked Ignis Elite holders.

The proof of concept is validated. The next phase is scaling capital, adding cbBTC trading pairs, and partnering with market makers to bring deep liquidity.

## 5. Economic Model & Fee Structure

### 5.1 Fee Architecture — Base (EVM)

Umbrae operates a dynamic fee model on Base through DLMM (Dynamic Liquidity Market Maker) pools. The protocol takes a percentage of the total dynamic fee on every swap before LP distribution.

**Protocol Take — Standard DLMM Pools:**

Cumulative Platform Volume	Protocol Take	LP Retention
\$0 (Launch)	5%	95%
\$500M	4.5%	95.5%
\$2B	4%	96%
\$5B	3.5%	96.5%
\$10B	3%	97%

Reductions are permanent and irreversible — enforced at the contract level. 3% is the floor.

**Platform Service Fees — Base:**

Service	Fee	Notes
Basic Swap (own pools)	0%	Revenue from protocol skim only
Routed Swap (external DEXs)	0.05–0.10%	Trades routed to Uniswap, Aerodrome, etc.
MEV Protection	0.08–0.15%	Private transaction routing
Automated Rebalancing	0.15–0.25%/month	Layer 2 gated feature
Copy Trading	0.08–0.10%/trade	Layer 2 gated feature

**MaaS Pools (Verified Projects):**

Pool Type	Protocol Take
MaaS Launch	15%
MaaS 6-month sustained	12%
MaaS 12-month sustained	10%

**Zero-Fee Operations (all chains):** LP entry/exit/repositioning, market analysis, portfolio tracking, price monitoring and alerts, custom settings and configuration.

For the complete smart contract implementation reference, including whitelisted pool configurations, anti-griefing defenses, and fee collection flow, see the *Umbræ Fee Infrastructure Spec v2.0*.

## 5.2 Revenue Distribution

We operate on a transparent distribution model. Every dollar earned in platform fees (Swaps, SAGE Credits, MM Services) is split as follows:

Allocation	%	Purpose
<b>Community</b>	<b>40%</b>	<b>Total returned to holders</b>
— OG Tier Holders	30%	Tiered yield to qualified NFT holders (Triple OG / Double OG / OG)
— Ignis Elite Holder Pool	5%	Direct yield to all Ignis Elite NFT holders — no OG tier required
— \$U1 Ecosystem	5%	Distributed to participating token holders
<b>Treasury</b>	<b>30%</b>	Bounties, future audits, legal defense, overhead
<b>Operations</b>	<b>30%</b>	Development, SAGE compute costs, salaries, infrastructure

The Ignis Elite Holder Pool was created when a former team member was removed. The 10% previously allocated to that role was restructured: 5% was absorbed by leadership to cover the additional workload, and the remaining 5% was proposed as a new community allocation. The proposal was presented to Ignis Elite holders, received a clear majority in favor, and was implemented. No existing allocation was reduced or diluted.

## 6. Tokenomics (\$U1)

The \$U1 token is not a “farm token”; it is a value-capture mechanism designed to permanently deepen ecosystem liquidity.

- ▲ **Total Supply:** 5,000,000 (Fixed, non-inflationary).
- ▲ **Circulating:** 4,000,000 (80%).
- ▲ **Protocol LP:** 1,000,000 (20% — Two-tier permanent liquidity model on Base).
- ▲ **Trading Pair:** \$U1/WETH on Base.
- ▲ **Base Fee:** 1% + dynamic variable fee (capped at 5%), down from 4% on the former Solana pool.

**Why ETH?** USDC has issuer-level freeze authority — Circle can blacklist any address holding USDC at any time. ETH has no issuer, no freeze function, and no centralized entity that can interfere with our liquidity. Pairing against ETH is the sovereign choice, consistent with the principles this platform is built on.

### 6.1 Token Migration: Solana → Base

\$U1 is being migrated from Solana to Base as an ERC-20. The tokenomics are preserved — this is a chain migration, not a redesign. Umbrae's entire platform infrastructure runs on EVM, and consolidating the token onto the same chain eliminates cross-chain friction and third-party dependency.

- ▲ **Conversion:** 1:1. Every Solana \$U1 converts to exactly 1 Base \$U1. No ratio changes, no dilution, no penalties.
- ▲ **Migration Window:** Open until November 2026. After that, unconverted Solana \$U1 is no longer eligible for migration. Unmigrated tokens are transferred to the Reserve LP Contract — they can only become permanently locked liquidity.
- ▲ **Contract Development & Security Review:** April 2026.
- ▲ **Pool Genesis:** Late April 2026 — \$U1/WETH pool deployed on Base.
- ▲ **Migration Opens:** Late April–Early May 2026.

For full details on the migration process, liquidity funding sources, timeline, and holder instructions, see the *\$U1 Token Migration: Solana → Base — Official Announcement* (available at [umbrae.io](https://umbrae.io)).

### 6.2 Protocol-Owned Liquidity: Two-Tier LP Model

20% of total supply (1,000,000 \$U1) is allocated as permanent protocol liquidity, deployed using a two-tier model designed to eliminate any possibility of a team dump.

**Tier 1 — Active LP:** 150,000 \$U1 paired against 30 ETH (~\$60,000), deployed directly into the \$U1/WETH DAMM pool. The LP tokens are permanently locked — they can never be withdrawn by anyone. This liquidity earns trading fees on every swap, claimed by the protocol for operational use.

**Tier 2 — Reserve LP:** The remaining 850,000 \$U1 sits in a smart contract with a single function: deposit into Active LP. These tokens cannot be withdrawn, sold, or transferred. The only direction they can move is into the permanently locked pool, paired against ETH. Every time the protocol commits more ETH, a proportional amount of Reserve \$U1 activates and joins the locked LP.

The pool can only grow. It can never shrink. This is verifiable on-chain. The Reserve LP Contract has no withdrawal function.

### 6.3 \$U1 Locking & Rewards

Once migrated, holders can immediately lock their \$U1 on the Umbrae platform. Locking does not require trading to be open — the platform is already live and generating fees.

**The Formula:**

Your weekly reward = (your locked \$U1 / total locked \$U1) × 5% of weekly platform fees.

**The Rules:**

1. Lock your \$U1 on the Umbrae platform at any time.
2. Must be locked for the full 7-day period to be eligible for that week’s rewards.
3. Unlock mid-week = forfeit that week’s rewards entirely.
4. Rewards paid every week on a fixed day.
5. Lock anytime, sell anytime — only locked tokens earn.

Staking and LP Provision will be deployed in subsequent phases after locking infrastructure is stable.

## 6.4 The “Compounding Liquidity” Flywheel

We solve the problem of mercenary capital via a mechanical buy-back loop:

1. **Fee Collection:** 5% of all Platform Fees are routed to the \$U1 Ecosystem pool.
2. **Distribution:** Yield is distributed weekly, proportionally to locked token holders.
3. **Protocol LP Growth:** As protocol tokens unlock from the Solana vesting contract (June–November 2026), they are converted to ETH and deposited into the permanently locked Base LP.

The result: liquidity depth increases over time, the price floor is mechanically raised, and lower slippage attracts larger volume — generating more fees and restarting the cycle.

## 7. Ignis Elite NFT System (Equity)

Access to the highest tier of rewards and Layer 2 (Gated) features is controlled by 777 fixed NFT licenses. These act as “Seats” on the exchange.

The purpose of Ignis Elite extends beyond yield and platform access. The entire focus of this system is to lay the foundations — through each holder — for decentralized sovereignty to exist in practice. Every Elite holder is a vetted participant in a network designed to operate independently of any centralized authority. The seats are limited because the integrity of the network depends on the quality of who holds them, not the quantity.

### 7.1 Tier Structure & Revenue Share

As outlined in Section 5.2, 30% of all platform fees go to the OG Tier pool. This 30% is distributed among qualified holders based on their tier:

Tier	Composition	Share of OG Tier Pool
Triple OG	Phase 1 + 2 + 3 (All in one wallet)	50% of the 30% OG pool
Double OG	Phase 1 OR (Phase 2 + 3× Phase 3)	30% of the 30% OG pool
OG	Phase 2 OR (3× Phase 3)	20% of the 30% OG pool

**Example:** If the platform generates \$100,000 in fees, \$30,000 goes to the OG Tier Pool. Triple OG holders split \$15,000 (50%), Double OG holders split \$9,000 (30%), and OG holders split \$6,000 (20%).

In addition, a separate 5% Ignis Elite Holder Pool is distributed to all Ignis Elite NFT holders regardless of OG tier. Hold any Ignis Elite NFT and earn from platform fees — no tier qualification required.

Holders can “Boost” their OG yield share by holding additional Phase 3 NFTs. There is a maximum boost cap per Ethereum address — excess Phase 3 NFTs beyond the cap do not provide additional boost credit. Refer to the boost formula and tier-specific caps on the Umbrae platform or [ignislabs.ai](https://ignislabs.ai) for current details.

## 7.2 Current Inventory & Pricing

Supply is strictly limited. Once sold out, entry is only possible via secondary markets.

- ▲ **Phase 1:** SOLD OUT.
- ▲ **Phase 2:** SOLD OUT.
- ▲ **Phase 3:** Still available for sale. Price: 1.0 ETH. Utility: Entry-level access; used as “Fuel” to boost yield for higher tiers. Check [umbrae.io](https://umbrae.io) or [ignislabs.ai](https://ignislabs.ai) for current remaining supply.

## 7.3 NFT Locking

Ignis Elite NFT locking goes live alongside the \$U1 token migration. Holders can lock their NFTs on the Umbrae platform to qualify for distributions from the Ignis Elite Holder Pool and arbitrage bot profits. Ignis Elite NFTs remain on Ethereum and are not affected by the \$U1 migration.

# 8. Token Verification & Market Making as a Service

Umbrae filters the noise of the crypto market through a strict curation process and offers professional liquidity services to verified projects.

## 8.1 Verification Tiers (Anti-Scam)

Tier	Designation	Meaning
Unverified	—	Token exists on-chain but has not been reviewed. “Trade at your own risk.”
Reviewed (Tier 2)	✓ Reviewed	Passed automated screening — no malicious code detected, contract safety checks cleared.
Verified Utility (Tier 1)	◆ Verified	Passed all three verification stages — automated screening, community review with council approval, and security audit.
Partner	◆ Partner	Deep integration with Umbrae infrastructure. Ongoing active collaboration.

## 8.2 The Three-Stage Verification Pipeline

- ▲ **Stage 1 — Automated Screening (AI + On-Chain Analysis):** Contract analysis, deployer wallet history, tokenomics assessment. Filters out scams before any human time is spent.
- ▲ **Stage 2 — Community Review + Council Approval:** Human evaluation by the Triple OG Council.
- ▲ **Stage 3 — Security Audit:** Independent verification of contract safety and team deliverables.

### 8.3 Market Making as a Service (MaaS)

MaaS is available exclusively to projects that have achieved Verified Utility (Tier 1) or Partner designation. MaaS includes curated DLMM pool configuration, professional liquidity bootstrapping, routing priority, SAGE analytics integration, and ongoing liquidity management support.

MaaS pools operate under a 15% protocol take at launch (vs. 5% standard), with milestone reductions to 12% at 6 months and 10% at 12 months for sustained performance.

For the full verification pipeline specification, see the *Umbrae Token Verification Pipeline* (separate document).

## 9. Gamification

### 9.1 The “RPG” of Trading

Umbrae transforms trading from a chore into a progression system.

- ▲ **Proficiency Ranks:** Users progress from Rank E to Rank S based on volume, profit ratio, and consistency.
- ▲ **Achievements:** Unlock badges for milestones (e.g., “DLMM Mastery”, “First \$1M Volume”).
- ▲ **Rewards:** Higher ranks unlock lower fees, priority support, and exclusive SAGE credits.

## 10. Roadmap: Milestone-Driven Execution

Umbrae’s roadmap is organized by milestones rather than hard calendar deadlines. Each phase is gated by the completion of its predecessor — we ship when it’s ready, not when a date arrives.

### 10.1 Completed

Milestone	Status
Base DLMM — First DLMM deployed on Base chain	✔ Complete
SAGE v1 — AI Credit System	✔ Complete
Token Verification Pipeline	✔ Complete
MEV Protection	✔ Complete
Proprietary Arbitrage Bot — Live on Base mainnet	✔ Complete
Hashlock Security Audit (Initial)	✔ Complete
Hashlock Security Audit (Follow-Up)	✔ Complete
\$U1 Token Launch	✔ Complete
Phase 1 NFT Sale	✔ Sold Out
Phase 2 NFT Sale	✔ Sold Out

## 10.2 Active & Upcoming

Milestone	Sequence	Gate
\$U1 Token Migration (Solana → Base)	Planned — Late April 2026	—
Ignis Elite NFT Locking	Launches with migration	—
Arbitrage Bot — Scale capital + cbBTC pairs	Active	—
\$U1 Locking Infrastructure	Next	Token migration underway
Private RPC Node Deployment	In Progress	—
DLMM Activation (post-migration)	After Migration	Token migration complete
\$U1 Staking Infrastructure	After Lock	Locking infrastructure live + stable
\$U1 LP Provision Infrastructure	After Stake	Staking infrastructure live + stable
Automated Rebalancing (Layer 2)	After Core Infra	KYC-gated, AI-managed LP strategies
Multi-Chain EVM Expansion (Arbitrum, Optimism, Ethereum, BNB, Avalanche, Hyperliquid)	After Core Infra	Core infrastructure stable, RPC nodes operational
Solana DLMM Deployment	After EVM	Base volume established, EVM expansion underway
SAGE v2 — Proprietary Model Training	After Data Threshold	Sufficient platform-native data for meaningful model advantage
Private Infrastructure Buildout (Full)	Ongoing	Revenue-funded, phased deployment

## 10.3 Execution Philosophy

We do not set public deadlines. The crypto space is littered with projects that over-promised timelines and shipped broken products. Our commitment is to milestone quality, not calendar performance. Each milestone listed above has clear completion criteria, and we communicate progress transparently to our community as work advances.

## 11. Team & Philosophy

Ignis AI Labs operates on a “Special Forces” philosophy: small, highly efficient teams rather than bloated headcounts.

### 11.1 Leadership

#### Elijah Matthew Moses — Founder & Chief Executive Officer

From the projects of Queens to breakthrough AI research. In crypto since 2012 — mining Bitcoin on high school computers when it was just a curiosity, running P2P crypto deals when there were no exchanges. Building code to survive, building platforms to thrive. Philosophy: protect the sovereignty of consciousness. The Emanon Principle — currently in development — will change how we understand intelligence itself. No VCs, no traditional funding — we build communities that understand true value. Serial builder from Silk Road-era crypto to developing MSRA, a novel AI architecture with 12.5M+ parameters — still in active development and already

demonstrating strong results on complex tasks. Hands-on experience in hardware architecture, data center buildouts, and trading infrastructure.

### **Ryan Stickney — Chief Operating Officer & MT5 Lead Developer**

First Triple OG holder and first major investor in Ignis — a journey from believer to builder. 27+ years of operational excellence across multiple industries, bringing enterprise-scale thinking to bleeding-edge innovation. Architecting the MT5 backtesting infrastructure and trading bot systems. Multi-founder currently running 2–3 companies alongside Ignis, including Benni Botanicals.

### **Gowtham Naidu Ponnana — Chief Security Officer**

At 20 years old, ranked #38 globally on HackenProof and #420 on Immunefi. \$560M+ in identified vulnerabilities across 60+ platforms; zero high-severity findings in Ignis systems. Former Head of Security at CryptoForce and PayOnRamp (India's fastest-growing crypto trading and on/off ramp platforms). DEFCON Delhi speaker; published research on ERC4626 inflation attacks and blockchain bridge vulnerabilities. Security partnerships with TECHFUND Inc (Japanese startup accelerator and auditing platform).

## **11.2 Philosophy**

- ▲ **Hiring:** We only hire individuals who can operate at “Founder Speed.”
- ▲ **Sustainability:** We are not reliant on VC funding. The platform is supported by diversified revenue streams (Trading Fees, Token Tax, Market Making Services).
- ▲ **Vertical Integration:** Unlike competitors who outsource infrastructure to cloud providers, our leadership can design and build hardware systems in-house. This enables the private infrastructure roadmap outlined in Section 4.3.
- ▲ **Security-First:** With a dedicated CSO embedded from day one, security is architected into our systems — not bolted on after the fact.

---

## **12. Disclaimer & Risk Factors**

### **Hybrid Compliance:**

Features requiring active management or specific regulatory oversight (Layer 2) are strictly gated via KYC. Umbrae is committed to operating within legal frameworks to ensure long-term viability.

### **Risk Warning:**

Trading crypto-assets involves significant risk. The “Compounding Liquidity” and “Revenue Share” models are subject to market volume and performance. All projections are forward-looking estimates and do not constitute guarantees of return. Ignis AI Labs does not provide investment advice.

---

**Document Version:** 4.0

**Confidential and Proprietary**

**Ignis AI Labs — March 2026**



Confidential and Proprietary