LITEPAPER

MIKO Protocol

An AI-Curated Solana Index of Weekly Assets

MIKO Protocol

Solana-Native Autonomous Agent and On-Chain Allocation Protocol

mikoprotocol.com · docs.mikoprotocol.com · June 2026

Abstract—The artificial-intelligence agent narrative captured 22.39% of crypto investor mindshare in 2025, yet AI-related tokens returned an average of −50.18% over the same period, ranking among the worst-performing categories in the market. The structural cause is a broken economic link: platforms capture protocol revenue, agents accumulate value in their own treasuries, and vault depositors earn yield, while the holder of an agent's token is left with pure price speculation. MIKO Protocol closes this gap. It is a Solana-native token ecosystem in which an AI agent's analytical output is structurally converted, every week, into on-chain asset selection, acquisition, and pro-rata allocation to eligible holders. A permanent 6% Token-2022 transfer fee funds an acquisition treasury; a four-layer intelligence stack — semantic ingestion, a knowledge graph with a six-provider fact-checking pipeline, a persona-driven generation core, and a self-improving three-phase machine-learning selection algorithm — chooses each week's asset; and an on-chain execution module harvests, swaps, and allocates the result. Every selection is recorded with its 24-hour and 7-day outcome, producing a publicly auditable track record. This litepaper summarizes the thesis, intelligence architecture, on-chain implementation, tokenomics, and risk model.

Index Terms—*autonomous AI agents, Solana, Token-2022, on-chain asset allocation, machine-learning ranking, multi-source fact-checking, tokenomics.*

I. Introduction: The Holder Outcome Gap

The convergence of artificial intelligence and cryptocurrency is the defining trend of the present market cycle. Capability has advanced rapidly: agents now hold wallets, execute decentralized-exchange (DEX) trades, manage treasuries, and deploy contracts autonomously. Investor attention has followed. According to the CoinGecko Annual Crypto Industry Report 2025, the AI-agent narrative captured 22.39% of investor mindshare in 2025, among the highest of any single category.

Returns, however, diverged sharply from attention. Over the same period AI-related tokens averaged a −50.18% return, one of the worst outcomes of any major category, despite being among the most popular. The reversal from the prior year is stark (Table I).

Table I: Mindshare versus average return for the AI-agent token category (CoinGecko, 2025).

Year	Mindshare	Avg. return
2024	15.67%	+2,940%
2025	22.39%	−50.18%

The pattern across the market is consistent. Agent capability is advancing, but the economic link between agent activity and what the token holder receives is broken in most projects. Platforms capture value through protocol revenue. Agents accumulate value in their own treasuries. Vault depositors earn yield. The holder of an agent's token, with no structured claim on the agent's activity, is left with pure price speculation. We refer to this as the **holder outcome gap**.

MIKO Protocol was built to close this gap. It rests on a single thesis:

> *An AI agent's analytical output drives weekly on-chain asset selection, acquisition, and allocation to its token holders, with a publicly verifiable track record.*

The remainder of this paper is organized as follows. Section II situates MIKO within the AI-agent market and isolates the unsolved problem. Section III states the thesis and gives a system overview. Section IV details the four-layer intelligence stack. Section V describes the on-chain implementation. Section VI presents tokenomics and distribution. Section VII addresses risk, security, and transparency. Section VIII outlines the roadmap, and Section IX concludes.

II. Background and the Unsolved Problem

The crypto AI-agent narrative began in 2024, when an autonomous agent promoting a memecoin demonstrated that a single AI posting on social media could move markets measured in billions. The year 2025 saw the buildout of infrastructure: Virtuals Protocol grew to more than 18,000 deployed agents and an Agentic GDP of $479M; elizaOS became the most-forked agent framework; and AIXBT reached 445,000+ followers by continuously monitoring 400+ key opinion leaders (KOLs). A general-purpose tool, OpenClaw, then democratized agent creation at a scale no crypto-native framework had achieved, reaching 234,000+ stars and spawning an ecosystem of derivatives.

By 2026 the landscape includes genuinely capable systems: regulated agentic wallets, the x402 protocol (115M+ machine-to-machine micropayments), protocol-level permission delegation (EIP-7702), and autonomous vaults managing tens of millions in total value locked (TVL). The agents of 2026 are not the chatbots of 2024.

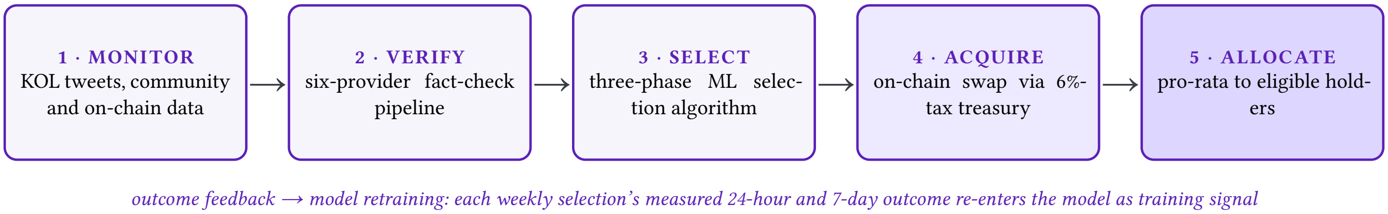

Fig. 1: The AI-to-allocation pipeline. The measured outcome of each weekly selection is scored and fed back as training signal, closing the loop between analysis and on-chain allocation.

Despite this progress, a structural gap persists across almost every project: the output of AI-agent activity does not reach token holders in a structured, automatic, and verifiable way. Table II summarizes the pattern.

Table II: Representative AI-agent projects: capability versus structured outcome to the individual token holder.

Project	Agent capability	Outcome to token holder
Virtuals	18,000+ agents; $479M aGDP	Platform captures revenue; no guaranteed per-agent holder return
AIXBT	445K followers; 400+ KOLs	Analysis to read; no on-chain delivery to wallets
Theoriq	$25M TVL autonomous vault	Returns accrue to depositors, not token holders
$CLAWD	Self-coding agent; treasury management	Agent retains treasury; no holder distribution

The conclusion is consistent: capability advances rapidly, while the economic link between agent performance and holder return remains broken or nonexistent. In a market where AI agents are becoming extraordinarily capable, the relevant question is *capable for whom.* If the answer is not "for the token holder," the capability is economically irrelevant to the person holding the token.

III. Thesis and System Overview

MIKO does not compete on agent autonomy, agent count, or general-purpose frameworks. It does one thing that almost no other project does: it takes the AI's analytical output, acquires the selected asset on-chain, and allocates it directly to every eligible token holder — every week, automatically, with a trackable performance record. The end-to-end pipeline comprises five stages (Fig. 1):

1. **Monitoring.** Miko continuously ingests hundreds of KOL tweets, community discussions, and on-chain market data across the Solana ecosystem.
2. **Verification.** A multi-source fact-checking engine verifies claims before they influence any decision, consulting up to six independent providers under an AI-driven verification strategy.
3. **Selection.** A self-improving machine-learning (ML) pipeline (Bayesian optimization ⟶ Thompson sampling ⟶ CatBoost learning-to-rank) selects the week's asset.
4. **Acquisition.** The on-chain execution module autonomously acquires the selected asset on a Solana DEX using accumulated transfer-tax revenue.
5. **Allocation.** The acquired asset is allocated to all eligible holders pro-rata to their $MIKO holdings.

A. Core Innovations

- **AI-to-allocation pipeline.** An AI agent's analytical output is directly converted into on-chain acquisitions and allocations to holders. Miko does not recommend; it acts.
- **Multi-source fact verification.** Before any information influences a weekly selection, it passes through an adaptive pipeline that consults multiple independent sources and requires evidence convergence.
- **Self-improving selection.** The selection algorithm evolves through three statistical phases, transitioning automatically as data accumulates and rolling back if performance degrades.
- **Sustainable on-chain funding.** A permanent 6% transfer tax, implemented via Solana's Token-2022 standard, provides a continuous, immutable funding stream. As long as $MIKO is traded, the acquisition treasury is funded.

B. A Distinct Category: AI-Curated Solana Index

Earlier Solana protocols (e.g., PRINT, IMG) established a tax-funded reward category that delivered a fixed reward asset, typically SOL. The static-asset model proved fragile: PRINT abandoned its reward model and IMG's market capitalization declined significantly from its peak. A fixed reward cannot keep up with a market defined by constant narrative rotation. MIKO operates in a different category. Each week the agent selects an asset from across the Solana ecosystem, the protocol acquires it via the accumulated transfer-tax treasury, and the position is allocated pro-rata. When the market rotates from DeFi to memecoins to infrastructure, MIKO's selection rotates with it. The force that erodes static models becomes the source of MIKO's value.

IV. The Intelligence Stack

Miko's intelligence is not a single model responding to prompts. It is a production system of specialized layers, each engineered for a specific function in the pipeline from raw data to on-chain execution.

A. Layer 1: Data Ingestion and Semantic Filtering

The first layer is the system's sensory apparatus, connecting to social and on-chain data sources in real time. Raw data is not useful data; this layer applies natural-language processing and heuristic rules to (i) filter noise, (ii) score relevance, (iii) classify content by topic and sentiment, and (iv) cluster related discussions using HDBSCAN over TF-IDF vectors and embedding similarity. Clustering is critical downstream: when several independent KOLs make the same claim, it carries different weight than a single uncorroborated source. The source weight of a claim c is

$$\mathrm{SourceWeight}(c) = \log\left(1 + \frac{n_{\mathrm{authors}}}{\alpha}\right) \times \overline{\mathrm{persuasion}} \quad (1)$$

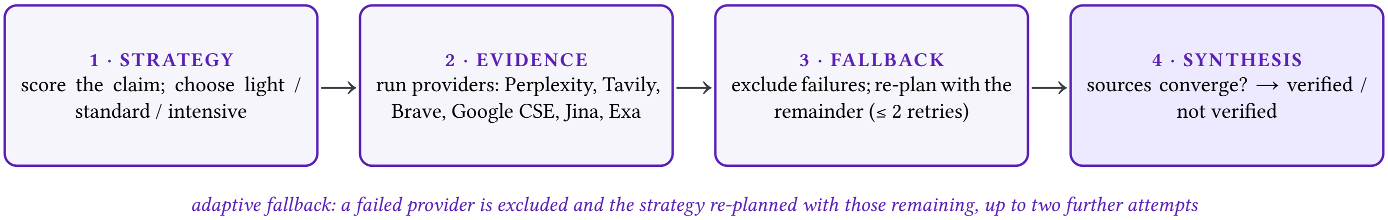

Fig. 2: The four-stage adaptive fact-checking pipeline. Verification intensity and provider choice are selected per claim; a claim is verified only when independent sources converge.

where n_{authors} is the count of independent sources making the claim and $\alpha = 0.8$ is a power-law exponent derived from the heavy-tailed distribution of social-media attention (Newman, 2005).

B. Layer 2: Knowledge Graph and Multi-Source Fact-Checking

This layer most fundamentally differentiates Miko from other agents. Information is not treated in isolation; it is integrated into a structured knowledge graph stored in PostgreSQL with vector embeddings for semantic search, mapping relationships among tokens, KOLs, market events, and projects. On encountering new information, Miko retrieves relevant context — past performance, associated KOLs, historical claims, verification status — enabling informed rather than reactive decisions. Semantic relevance uses cosine similarity on embedding vectors,

$$\text{relevance}(q, d) = \frac{\vec{q} \cdot \vec{d}}{\|\vec{q}\| \, \|\vec{d}\|} \tag{2}$$

with results above an importance threshold ($\tau = 0.4$ by default) surfaced as contextual grounding for both content generation and asset selection.

The fact-checking engine is a four-stage adaptive pipeline consulting six independent verification providers (Perplexity, Tavily, Brave Search, Google CSE, Jina, and Exa), as shown in Fig. 2. Its key design decisions are:

- **AI-driven strategy selection.** The verification level is not hardcoded. The AI evaluates each claim's specificity, importance, and independent-source count, then chooses the intensity (light: one provider; standard: two; intensive: three or more) and the specific providers whose strengths match the claim type.
- **Adaptive fallback.** If a provider fails, the system excludes it and re-plans with the remainder, for up to two additional attempts. Temporary outages never block verification.
- **Evidence convergence.** A claim is verified only when evidence from independent sources converges on the same conclusion. Divergent or insufficient evidence yields a "not verified" verdict.
- **Specificity-aware evaluation.** Synthesis checks whether specific details (names, numbers, dates, amounts) match the evidence, not merely the general topic.

The pipeline operates at three independent checkpoints: proactive content (a topic is blocked from publication unless verified), mention responses (verification status sets the confidence of replies), and asset selection (verification status weights candidate scoring). For holders, this means information feeding a selection has been subjected to multi-source verification — the layer standing between a holder's weekly allocation and a decision based on false information.

C. Layer 3: Persona-Driven Generation Core

This layer creates Miko's public-facing voice on social media, which drives community engagement and growth. It runs on a multi-model architecture (Gemini and Claude) with automatic failover, selects among multiple reaction modes by evaluating conversational context, and verifies every response for originality before posting. The economic relevance is direct: social presence drives community growth, which drives trading volume, which generates transfer-tax revenue, which funds the weekly acquisitions. The quality of Miko's social output thus influences the size of the acquisition treasury.

D. Layer 4: Selection Algorithm

This is the financial brain of the protocol: the system that determines which asset holders receive each week. It evolves through three statistical phases, each requiring more data and delivering higher precision; transitions are automatic (on statistical-significance thresholds) and reversible (automatic rollback on degradation), as summarized in Table III.

Table III: Three-phase evolution of the selection algorithm. ρ is the Spearman correlation between predicted and realized outcomes.

Phase	Method	Min. data	Conf.	Threshold
1 — Exploration	Bayesian opt. (Optuna)	6 wks	80%	$\rho \geq 0.65$
2 — Exploitation	Thompson + ε-greedy	16 wks	85%	$\rho \geq 0.70$
3 — Precision	CatBoost LTR (YetiRank)	26 wks	90%	$\rho \geq 0.75$

In Phase 1, with limited history, Bayesian optimization explores which features and weight combinations best predict successful selections. In Phase 2, Thompson sampling with 10% epsilon-greedy exploration ($\varepsilon = 0.10$) balances exploitation of known patterns with continued exploration. In Phase 3, a CatBoost model with the YetiRank loss predicts which candidate will achieve the highest composite outcome score. Each selection is evaluated against a multi-metric composite,

$$S_{\text{composite}} = \sum_{i=1}^{n} w_i x_i, \quad \sum_{i=1}^{6} w_i = 1.00 \tag{3}$$

whose weights are derived from empirical research on cryptocurrency success factors (Table IV).

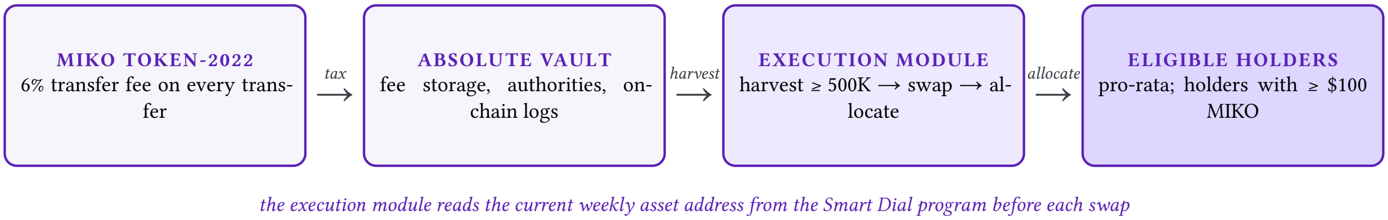

Fig. 3: On-chain architecture. Verified programs (Token-2022, Absolute Vault, Smart Dial) enforce the rules; the execution module harvests fees, swaps for the selected asset, and allocates pro-rata.

Table IV: Composite outcome score: metrics, weights, and research basis.

Metric (x_i)	w_i	Research basis
Selected-asset price performance	0.40	Chen et al. (2023)
$MIKO holder growth rate	0.22	Liu and Tsyvinski (2021)
$MIKO token price performance	0.20	Eisenmann et al. (2006)
Community sentiment	0.10	Kraaijeveld and De Smedt (2020)
Tweet engagement	0.05	DeFi Alliance (2024)
$MIKO volume growth	0.03	Brandvold et al. (2015)

Hard-coded safeguards revert the model to its previous phase when performance degrades. Rollback triggers if performance drops at least 15% from baseline (Δ_{perf}), three consecutive selections underperform (n_{fail}), or model confidence degrades by 10% (Δ_{conf}):

$$\Delta_{\text{perf}} \geq 0.15, \quad n_{\text{fail}} \geq 3, \quad \Delta_{\text{conf}} \geq 0.10 \tag{4}$$

This ensures a poorly calibrated model never persists in committing real capital.

1) Token Quality Filters

Before consideration, a candidate must pass a two-tier assessment. Tier 1 applies hard thresholds (Table V). Tier 2 then evaluates surviving candidates through a multi-factor DEX market-structure analysis — order flow (buy/sell pressure), transaction velocity, breakout readiness, relative strength, and holder breadth — producing a market bonus subject to multiplicative safety penalties (e.g., 75% for severe risk flags such as concentrated ownership; 50% for fast-decay patterns; a proportional discount for elevated risk).

Table V: Tier-1 hard-threshold quality filters.

Filter	Threshold	Rationale
Min. market cap	$3,000,000	Manipulation risk below (Kaiko, 2025)
Min. 24h volume	$1,000,000	Liquidity for execution
Excluded	$MIKO	Prevents self-selection conflict
Exempt	$SOL	Native asset exemption

2) Community Suggestions and Persuasion

Community members can recommend tokens by mentioning the agent with a $SYMBOL tag. Unlike vote-counting, each recommendation is scored by a persuasion function

$$\begin{aligned} \text{persuasion}(t) = f(&\text{authenticity}, \text{reasoning}, \\ &\text{alignment}, \text{relevance}) \end{aligned} \tag{5}$$

A minimum persuasion threshold of 0.3 (lower 30th-percentile cutoff) and a minimum of three total mentions ensure that only reasoned, authentic input influences the selection; repeated spamming of the same symbol is ineffective.

3) The Self-Improvement Loop

Every selection generates outcome data. An outcome collector measures the selected asset's 24-hour and 7-day price performance, computes the composite outcome score, appends features and outcome to the training dataset, and retrains the model for the next cycle. Early selections are exploratory; later selections are informed by months of accumulated outcomes and a fully trained ranking model. This is the source of MIKO's verifiable track record: every selection is recorded with the token chosen, the ML phase and model version, whether it was an exploration pick, and its realized outcome. Phase transitions and rollbacks are logged with their trigger metrics, producing a complete, queryable history.

V. On-Chain Implementation

Miko's intelligence stack produces analytical decisions off-chain; these are executed on Solana through verified on-chain programs and an on-chain execution module (Fig. 3). The hybrid architecture maximizes analytical flexibility while preserving on-chain transparency: every deployed program is verifiable against its published source code, so any change to on-chain logic is itself public. On-chain programs perform only the minimum necessary functions to reduce attack surface.

A. Core On-Chain Components

- **MIKO Token (Token-2022).** The core asset imposes a fixed 6% fee on all transfers (buys, sells, wallet-to-wallet) via the TransferFeeConfig extension. The mint authority is renounced and the freeze authority disabled, fixing total supply permanently and preventing wallet freezes.
- **Absolute Vault.** The protocol's central treasury. It manages transferred authorities, stores harvested transaction fees, and records activity on-chain for transparency.
- **Smart Dial.** A minimal configuration program storing the address of the current weekly selected asset, queried by the execution module during each allocation.

B. On-Chain Execution Module

The execution module converts analytical decisions into blockchain transactions, handling complex swap logic and periodic tasks off-chain while finalizing transactions on-chain. Its responsibilities are: **automatic harvesting** — when withheld fees reach 0.05% of total supply (500,000 MIKO), they are harvested into the Absolute Vault; **swap and allocation** — harvested MIKO is swapped for the AI-selected asset on a Solana DEX and allocated pro-rata, with a portion of collected tax directed to protocol operations (infrastructure, servers, development); **dynamic eligibility screening** — holder wallets are scanned in real time so that only holders with $MIKO value of $100 or more receive allocations; and **smart exclusion** — system-related addresses such as liquidity-pool and routing-contract accounts are detected and excluded so benefits reach actual community members. Each holder's weekly allocation is

$$\text{Allocation}_w = \frac{\text{HolderBalance}}{\text{TotalEligibleSupply}} \times \text{Treasury}_w \quad (6)$$

There is no platform intermediary, staking requirement, or governance vote: hold the token, receive the weekly allocation. The on-chain programs enforce the rules and protect assets; the execution module acts autonomously within those rules, completing the pipeline from analysis to allocation without manual intervention.

VI. Tokenomics and Distribution

MIKO's tokenomics prioritize fairness, long-term sustainability, and community trust, and are designed to mitigate the risks associated with team token holdings.

A. Specifications

Token name: MIKO Protocol. Ticker: $MIKO. Blockchain: Solana (SPL Token-2022 standard). Total supply: 1,000,000,000 MIKO. The contract address is published at deployment.

B. Initial Distribution

MIKO is committed to a 100% fair launch with no presales, private sales, or bundled launches, fostering an organic community from day one. No one, including the team, acquires tokens below the market price, and every allocation outside the liquidity pool starts inside an on-chain vesting escrow. The allocation is summarized in Table VI.

Table VI: Initial token distribution (total supply 1,000,000,000 MIKO).

Allocation	Share	Purpose
Liquidity pool	55%	Deposited at launch; LP position locked for 12 months
Community	13%	Airdrop seasons and events; 0.3% of supply unlocks each Tuesday at 00:00 UTC
Growth and partnerships	10%	Marketing and KOL collaborations; vests 1% of supply per 30 days
Ecosystem reserve	10%	Operations, liquidity injections, backstop; vests 2% of supply per 30 days
CEX listings	7%	Exchange listing inventory; 3-month milestone lock, re-locked if unused
Migration reserve	5%	Subsidizes the 6% fee for exchange-to-self-custody withdrawals; 3-month lock

C. Locked Allocations and Team Exposure

Every allocation outside the liquidity pool is held in its own vesting escrow on Jupiter Lock, an audited open-source locker with cancellation and recipient changes permanently disabled, and each escrow address is published for on-chain verification. The schedule itself caps the team's unlocked, spendable holdings at about 3.3% of total supply at any time: the current ecosystem-reserve tranche (2% of supply per 30 days across five tranches), the current growth tranche (1% of supply per 30 days across ten tranches), and the current weekly community tranche (0.3% of supply, unlocking each Tuesday at 00:00 UTC). The CEX listings inventory (7%) and the on-chain migration reserve (5%, which subsidizes the 6% transfer fee for users withdrawing from exchanges into self-custody) unlock only at a 3-month growth milestone and are locked again if unused. At the end of each vesting cycle, any unspent balance in the team wallet is burned or added to protocol liquidity, with both actions verifiable on-chain. These allocations fuel the project's early growth and, as the protocol matures, return any surplus to ecosystem health through a permanently smaller supply or deeper liquidity.

VII. Risk, Security, and Transparency

Because MIKO combines autonomous systems with financial incentives, it employs multi-layered measures to protect users and maintain trust.

A. Contract Security and Immutability

The token has no backdoors: the mint authority is permanently renounced and the freeze authority disabled, so total supply cannot change and no wallet can be frozen. Off-chain execution is permissioned but auditable: all fund movements require the execution module's operational key signature, are recorded in on-chain logs, and are verified by dedicated verification contracts that mechanically check core logic (e.g., 6% fee accuracy, system-account exclusion) at each deployment stage.

B. Fair-Distribution Safeguards

Initial liquidity-pool accounts and team- and development-related system accounts are excluded from allocation, so acquired assets reach only genuine community holders. Strict on-chain filtering (volume, liquidity, market cap, holder count) prevents scam or inactive tokens from being selected. As a final safeguard, in the rare event of a sudden serious issue with a selected asset (such as a rug-pull), the team holds limited authority to make an emergency change to the weekly selected asset, restricted to once per week.

C. Market Risk and Mitigation

The cryptocurrency market is highly volatile; the value of $MIKO and of any weekly selected asset can fluctuate sharply, and there is no guarantee that selections lead to profit. Mitigations are structural rather than promissory: weekly rotation of the selected asset diversifies holder exposure over time; the protocol's value is measured by long-term ability to generate alpha above the market average rather than by any single week's price action; and minimal management functions (limited to actions that do not harm the token's economic structure, such as a once-weekly emergency asset change or recovery from system error) are recorded on-chain so their use is restricted to protecting holder assets.

VIII. Roadmap

Development proceeds in phases, with the AI model continuously enhanced throughout.

- **Phase 1 — Genesis.** Launch the Miko AI agent (KOL ingestion, knowledge base, fact-checking, weekly selection); deploy the Token-2022 contract and allocation programs on mainnet; provide initial liquidity and begin trading; commence the weekly allocation cycle; and expose the intelligence stack through a wallet-authenticated REST API and MCP server.
- **Phase 2 — Ecosystem Tools.** Launch the MIKO's Insight dashboard and a public Selection Track Record; introduce

KOL Spotlight; and formalize partnerships with projects selected as weekly assets.
- **Phase 3 — Community Integration.** Launch the MIKO's Circle platform linking social and on-chain activity; introduce a participate-to-earn mechanism; and expand beyond X to additional social platforms.
- **Phase 4 — Intelligence Infrastructure.** Expose the fact-checking and market-analysis capabilities as an intelligence-as-a-service API for other autonomous agents, and operate an AI-curation launchpad whose revenue further funds the acquisition treasury.

IX. Conclusion: Intelligence That Pays

The AI-agent era has produced more agents, more autonomy, and more economic activity, and yet, for most projects, the link between that capability and the token holder's wallet does not exist. The gap is not inevitable; it is a design choice, and MIKO chose differently. The $MIKO token is the medium through which an AI's analytical output is converted into weekly on-chain asset allocations for its holders. Every week, the intelligence stack processes hundreds of data points, a multi-source pipeline verifies information before it influences decisions, a self-improving ML system selects the asset, the on-chain module acquires it with accumulated tax revenue, and every eligible holder receives the acquired asset pro-rata. The system is measurable by default: each selection is recorded with its 24-hour and 7-day outcome and a composite score that feeds back into learning. MIKO does not claim to be the most autonomous agent in the market. It does one thing: it converts AI analysis into a weekly on-chain allocation in the holder's wallet. That is the standard for an intelligent asset, and the blockchain keeps the score.

References

Brandvold, M., Molnár, P., Vagstad, K., and Valstad, Ø. C. A. (2015). Price discovery on Bitcoin exchanges. *Journal of International Financial Markets, Institutions and Money.*

Chen, et al. (2023). Empirical factors in cryptocurrency return performance.

CoinGecko. (2025). *Annual Crypto Industry Report 2025.*

DeFi Alliance. (2024). Engagement indicators in decentralized finance.

Eisenmann, T., Parker, G., and Van Alstyne, M. W. (2006). Strategies for Two-Sided Markets. *Harvard Business Review.*

Kaiko Research. (2025). Market structure and manipulation risk in low-capitalization tokens.

Kraaijeveld, O., and De Smedt, J. (2020). The predictive power of public Twitter sentiment for forecasting cryptocurrency prices. *Journal of International Financial Markets, Institutions and Money.*

Liu, Y., and Tsyvinski, A. (2021). Risks and Returns of Cryptocurrency. *The Review of Financial Studies.*

Newman, M. E. J. (2005). Power laws, Pareto distributions and Zipf's law. *Contemporary Physics.*

Reference list follows the sources cited in the MIKO Protocol technical documentation (whitepaper); bibliographic detail is given to the extent stated in the source material.