DAPPOS: The Web3 AI Operating System

DAPPOS Research

September 23, 2025

1 Introduction

As AI permeates every industry, generic, one-size-fits-all models are no longer sufficient for Web3, which demands a vertical AI with deep domain expertise, reasoning within specific constraints, and reliable execution of user intents. Prevailing general-purpose AI models, like those in ChatGPT or Perplexity, excel in broad tasks such as summarization or code generation but falter in Web3's dynamic ecosystem of DeFi, NFTs, meme coins, and blockchain due to deficiencies in Intelligence and Execution [1–6]. In Intelligence, they lack secure access to real-time, user-specific blockchain data—leading to imprecise recommendations for strategies like DeFi lending—while applying uniform priors that undervalue high-signal events (e.g., meme coin partnerships or token listings) and misinterpret Web3's nuances, lexicon, and risks like bot-driven hype or veiled indicators [1, 6]. In Execution, these models are limited to consultation, forcing users to manually handle wallets, DEXs, bridges, gas fees, slippage, and cross-chain complexities, hindering practical utility [10–13].

DAPPOS is a Web3 AI Operating System (OS) that intelligently researches, plans, and executes anything crypto-related. The Web3 AI OS is structured into two core layers: the Intelligence Layer and the Execution Layer.

The Intelligence Layer establishes the cognitive foundation, powered by a scalable Multi-Agent Framework (MAF) with 300–400 specialized vertical agents and more than 200 integrated tools [2–5]. This architecture enables broad Web3 generalization and adaptive intelligence, excelling across diverse domains—from quantitative analysis for trading strategies to marketing insights—while surpassing generalist AIs such as Perplexity in contextual reasoning and task-specific optimization.

At the core of the Intelligence Layer lies the Bubble Engine, driven by continuous reinforcement learning (RL) models with contextual RAG and compound memory tailored for Web3. It evolves perpetually by ingesting real-time insights from sources such as X and Binance Square, enabling rapid, iterative adaptation to new knowledge [1, 6–9]. Through the Bubble task platform, users can accelerate this process by posting tweets that "force-feed" the system novel concepts; once verified, these are seamlessly integrated, resolving challenges that may have seemed unsolvable just days earlier.

The Execution Layer complements the Intelligence Layer through DAPPOS's proven Intent Execution Network—a stable, time-tested infrastructure that delivers institutional-grade efficiency and safety [11–17]. Hardened by over 5 million users and 12 million transactions, it reliably transforms AI-generated plans into secure, on-chain interactions.

Built on the Web3 AI Operating System, DAPPOS introduces Instant dApps—enabling users to create and deploy Web3 applications on demand through simple natural language interactions [2–4]. The Discover Page acts as a dynamic hub, showcasing AI-generated insights and execution plans shared by the community. Users can browse, explore, and adopt strategies that align with their needs, executing them directly if desired. This experience mirrors vibe coding for websites or dashboards—where ideas transform seamlessly into deployable tools.

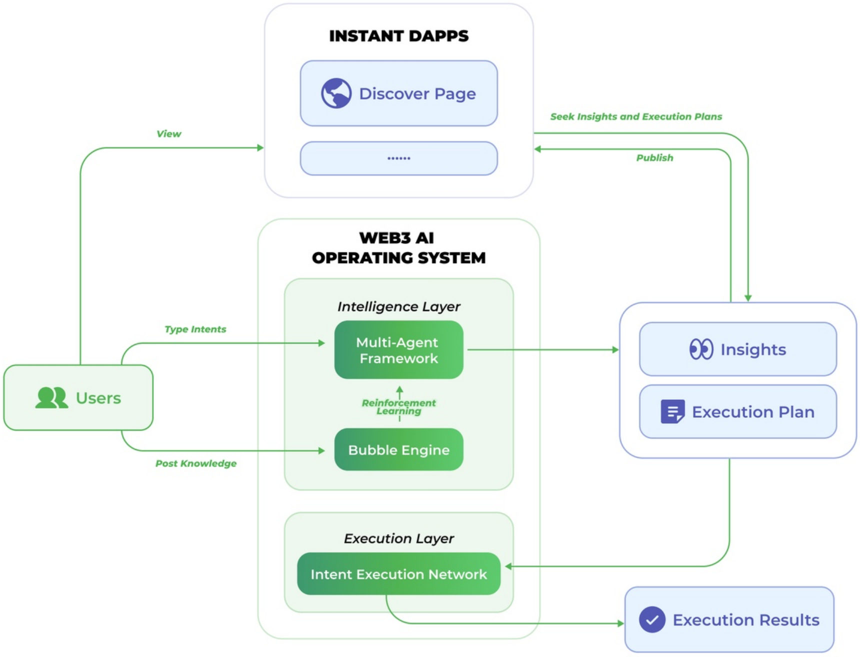

Figure 1: The User Journey of DAPPOS

Through the Web3 AI OS, DAPPOS unifies the fragmented Web3 landscape, enabling users to orchestrate complex actions through simple conversations as AI seamlessly coordinates across decentralized protocols. This one-stop system covers the entire Web3 lifecycle—from research and strategy design to autonomous execution—lowering barriers to entry and catalyzing innovation across the ecosystem. Ultimately, DAPPOS marks a paradigm shift for Web3, transforming it into an accessible frontier where anyone can engage, build, and create within the Web3 AI OS—unlocking the full potential of decentralized technologies for the advancement of human society.

2 The Intelligence Layer: DAPPOS AI

2.1 Core Components

The Intelligence Layer of DAPPOS forms the cognitive backbone of the Web3 AI Operating System, orchestrating research, planning, and decision-making for crypto-related tasks. It is structured of two core components: the Multi-Agent Framework (MAF) and the Bubble Engine, which together enable scalable, adaptive intelligence tailored to Web3's dynamic environment.

The Multi-Agent Framework (MAF) serves as the orchestration module, comprising 300–400 specialized vertical agents and over 200 integrated tools. Built on a stateful graph backbone, MAF facilitates composable autonomy through deterministic routing, subgraphs for tasks, and advanced features such as multi-agent search and tool-augmented generation [2–5]. This ensures explainable, production-grade AI that excels in domain-specific reasoning and generalization across quantitative analysis, trading strategies, and market insights.

Complementing MAF, the Bubble Engine acts as the adaptive learning core, powered by continuous reinforcement learning (RL) models optimized for Web3. It perpetually evolves by ingesting real-time data from sources like X and Binance Square, incorporating user-contributed insights via the Bubble task platform. Key mechanisms include Contextual Retrieval-Augmented Generation (RAG) for provenance-first data aggregation and Compound Memory for blending durable knowledge with episodic learnings, while handling misinformation and leveraging Web3 incentives for collaborative growth [1, 6–9].

Together, MAF and the Bubble Engine help DAPPOS create a synergistic system that outperforms generic AI in Web3, delivering excellent and evolving intelligence ready for the Execution Layer.

2.2 Multi-Agent Framework (MAF)

The Multi-Agent Framework (MAF) in DAPPOS represents a groundbreaking evolution in AI orchestration for Web3, guaranteeing explainable AI decisioning and production-grade reliability.

2.2.1 Stateful Graph Backbone: The Foundation of Composable Autonomy

At the core of MAF lies the StateGraph, a structured state schema composed of a collection of variables that forms the backbone for all operations. The StateGraph encapsulates key system elements such as messages for inter-agent communication, plans for structured workflows, artifacts (e.g., data outputs or models), UI interrupts for human oversight, final outputs, errors for robust handling, and flags for conditional logic. These variables are passed seamlessly across each state in the system, serving as the decision-making foundation for controller nodes at each stage [2, 4, 5].

In MAF, each node represents an AI agent for a designed task. Specifically, controller nodes evaluate the current state to determine the next flow direction and update the state variables accordingly. The process of a user query begins with the main controller task_planner_node, which analyzes user intents to classify the task type—such as information gathering, strategy execution, or alpha detection—and deterministically routes it to the appropriate subgraph (search, DeFi, or opportunity).

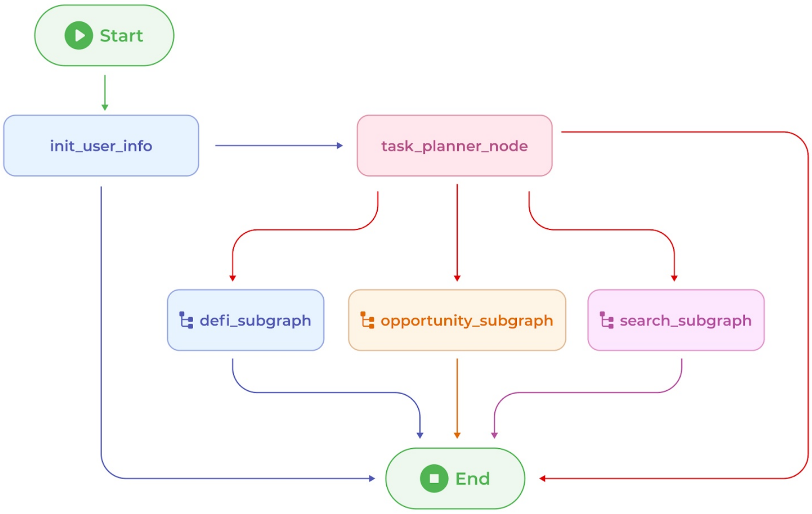

Figure 2: The Process of a User Query

Within each subgraph, specialized controller nodes handle internal flows, ensuring predictable paths that minimize variability and enhance trust in dynamic Web3 environments. This setup powers deterministic controllers and AI interrupt mechanisms, preventing errors while allowing freedom in AI decision-making.

2.2.2 Subgraphs: Specialized Modules for Web3 Intelligence

MAF's subgraphs are modular building blocks, each tailored to specific Web3 domains with deterministic node flows that process data efficiently. These subgraphs leverage the stateful backbone to integrate seamlessly, enabling adaptive intelligence across trading, analysis, and execution. For example:

Search Subgraph: This subgraph combines autonomous controllers with enhanced review cycles for continuous deepening and refinement, embodying an agentic search loop. The flow starts with search_controller_node that selects and parallelizes agents (e.g., google_search_agent_node for web data, twitter_search_agent_node for social sentiment, parallel_search_coordinator_node for multi-source queries). Outputs are merged in merge_results_node, filtered for relevance, and refined through an iterative enhanced review loop until confidence criteria (e.g., source diversity or accuracy scores) are met, which is checked by enhanced_search_review_node. The summarize_node then produces cited answers with frontend-friendly markers, ensuring transparency and usability [1–3, 7–9].

DeFi Subgraph: Centered on tool-augmented execution, this subgraph generates and validates DeFi strategies using over 200 integrated tools for precise operations like API queries or trade simulations. It begins with fetch_market_data to retrieve real-time blockchain metrics, followed by strategy_generator to create optimized

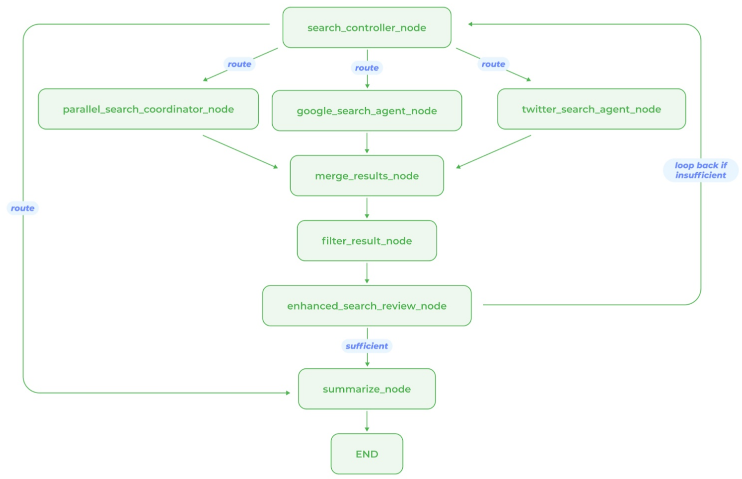

Figure 3: Search Subgraph in MAF

plans (e.g., yield farming or liquidity provision). The strategy_validator closes the loop by assessing risks, simulating outcomes, and incorporating optional user choices [2, 3].

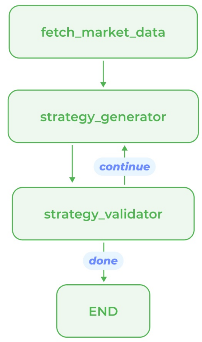

Figure 4: DeFi Subgraph in MAF

Opportunity Subgraph: This subgraph identifies high-potential Web3 opportunities by scanning markets, trends, and signals. It flows through opportunity_search_node for initial detection, fetch_market_data for validation, and an opportunity_node that synthesizes insights into actionable recommendations, such as emerging meme coins or protocol upgrades [2, 5].

2.2.3 Advanced Features: Enhancing Collaboration, Execution, and Interactivity

Building on the core subgraphs, MAF incorporates advanced capabilities to foster human–AI synergy and real-time adaptability. Multi-Agent Search: Central to the search subgraph, this feature enables autonomous and collaborative intelli-

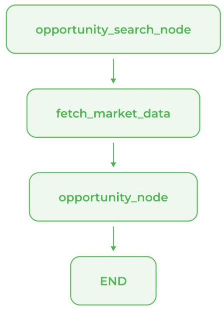

Figure 5: Opportunity Subgraph in MAF

gence, where a controller dynamically selects agents based on query complexity for parallel execution, boosting speed and breadth. For example, evaluating a meme coin's potential might invoke Google agents for news, Twitter agents for sentiment, and parallel custom agents for on-chain data. The enhanced review loop iterates refinements, and the summarizer outputs structured, cited responses optimized for frontend display, with HCI interrupts for real-time oversight in mission-critical actions [2–5]. Tool-Augmented Plan Generation: Particularly prominent in DeFi and opportunity subgraphs, this integrates structured tools to transform abstract insights into executable strategies. Validator nodes ensure safety by cross-checking against risk parameters, with optional UI interrupts for user approvals, closing the intelligence-action loop and preparing plans for seamless handover to the Execution Layer [2, 3]. Streaming and Interrupts: To support interactive experiences, MAF streams messages during long-running tasks, keeping users informed in real-time. When user queries or provided information are insufficient for accurate analysis, the AI pauses execution to seek clarification by prompting the user with targeted questions, ensuring completeness and reliability [6]. In essence, DAPPOS's Multi-Agent Framework redefines Web3 intelligence by combining composable autonomy with safety rails, enabling users to tackle complex tasks with unprecedented precision and adaptability. This framework not only outperforms generic AIs in domain-specific reasoning but also paves the way for collaborative, evolving ecosystems where AI and humans co-create value.

2.3 Bubble Engine: Continuous RL Improvement

At the core of DAPPOS's Intelligence Layer is the Bubble Engine, a Web3-focused reinforcement learning (RL) engine that drives perpetual evolution and adaptive intelligence. Powered by advanced mechanisms like Contextual Retrieval-Augmented Generation (RAG) and Compound Memory, the Bubble Engine transforms fragmented Web3 data into coherent, actionable intelligence.

2.3.1 Contextual RAG: Provenance-First Intelligence Aggregation

The Bubble Engine's Contextual RAG forms the backbone of its information processing, unifying diverse sources—web content, internal documents, and on/off-chain signals—into a provenance-first context layer that prioritizes traceability and reliability. This hybrid retriever combines dense (semantic) and sparse (keyword-based) techniques with domain-specific routing to orchestrate sub-queries efficiently. It deduplicates redundant results, reranks them using cross-encoders for relevance, and assembles a minimal sufficient evidence pack tailored to the query's needs [1, 7–9].

The generation process is citation-locked, producing 1]-style inline references with graceful fallbacks when evidence is limited. Dynamic budgets intelligently adjust retrieval depth and parallel agent execution based on detected uncertainty levels, optimizing for efficiency in volatile Web3 scenarios. Episodic memory captures short-term learnings from recent interactions, while long-term memory preserves confirmed facts for future recall [6]. Looking ahead, DAPPOS sets the stage for multimodal RAG (incorporating images, videos, and audio) and graph-based extensions, enabling richer contextual intelligence that connects entities, relationships, and temporal dynamics across the blockchain landscape [7–9].

2.3.2 Compound Memory: Durable and Episodic Knowledge Evolution

Complementing Contextual RAG, the Bubble Engine's Compound Memory integrates durable knowledge—core, verified Web3 facts and models—with episodic learnings accumulated over time. This approach ensures the engine retains timeless domain expertise (e.g., DeFi protocols and blockchain fundamentals) while dynamically incorporating transient insights, such as real-time market shifts or user-submitted alphas. As new knowledge is verified and integrated, the memory system evolves, creating a compounding effect that enhances reasoning accuracy and personalization [6].

2.3.3 Handling Misinformation and User Incentives

To maintain integrity in Web3's deception-prone environment, the Bubble Engine employs advanced fake information handling, cross-verifying sources against on-chain data and community flags to detect and mitigate bots, pumps, or manipulated signals [10, 18–20]. Additionally, Web3 incentives motivate users to share high-fidelity insights, accelerating collective intelligence growth.

In summary, the Bubble Engine redefines adaptive AI for Web3 by blending continuous RL with sophisticated retrieval and memory systems, delivering hyper-relevant, trustworthy insights that evolve in lockstep with the ecosystem.

3 The Execution Layer: Intent Execution Network

3.1 Current Intent-based Systems and Their Limitations

Intent, in the context of blockchain transactions, refers to the desired outcomes that users want to achieve, rather than specifying the exact steps needed to reach those outcomes. For instance, instead of detailing every interaction with smart contracts, nonce management, and gas payments, users can express their desired result in a straightforward manner, leaving the intricate execution details to sophisticated third parties [11]. This approach not only simplifies user interactions but also enhances efficiency by leveraging the expertise of third-party executors. Intents can range from basic needs to broader, more complex requirements. As long as users are willing to pay a certain cost and expect a specific result, intents can cover a wide variety of demands. For example:

- Buying BTC with 1000 USDT.
- Receiving a specified amount of USD in a certain bank account.
- Spending a certain amount of USDT or USD to receive the latest iPhone.

In conclusion, a common definition of intents is "I want X and I'm willing to pay up to C." Intents can be considered as signed messages that allow for a set of state transitions from a given starting state. Unlike traditional transactions that specify a unique computational path, intents allow for any path that satisfies certain constraints. This flexibility enables the matching of overlapping intents, increasing gas and economic efficiency [11, 14–17]. Current intent-based systems can be broadly classified into several categories based on their design and functionality:

- **Limit Orders and Auctions**: Systems like CowSwap and 1inch Fusion utilize limit orders and batch auctions to match intents. These systems focus on finding the best execution prices by promoting competition among liquidity providers. However, they are often limited to specific use cases and lack the flexibility to handle more complex intents [14–16].
- **Account Abstraction (AA)**: Proposals like EIP 4337 introduce account abstraction, allowing users to express intents at the wallet level. This approach decouples accounts from signers, enabling custom authorization logic tailored to users' needs. However, it introduces new complexities and requires substantial changes to Ethereum's core protocol [17, 21–23].
- **Generalized Intent Infrastructure**: Projects like Anoma and Essential aim to provide a comprehensive intent-based architecture. Anoma, for instance, utilizes a homogeneous protocol with a heterogeneous security model, allowing for decentralized counterparty discovery and solving. Essential, on the other hand, focuses on providing a modular intent layer and a new constraint-based language for more efficient intent executions. However, these projects are still under development and not accessible to users [11–13].

While intent-based designs offer significant improvements over traditional transaction methods, they are not without limitations. Current solutions often face several critical issues:

- **Limited Generalizability**: Due to the constraints of current validation mechanisms, intent-based systems often lack the flexibility to be applied across a wide range of decentralized applications (dApps). These systems are typically designed to solve specific problems and are not adaptable to various use cases. For example, a system optimized for limit orders may not be suitable for more complex intents such as cross-chain transactions or interactions with multiple DeFi protocols. The validation process often restricts the types of intents that can be executed, limiting the system's generalizability [11–13].

- **Restricted Solution Space**: Current solutions require synchronous execution, where all transaction components must be processed within a single blockchain block, and they lack mechanisms for pre-authorization, forcing users to manually approve each transaction. The inability to handle asynchronous operations and the lack of pre-authorization severely limit the solution space of existing intent-based systems. A practical example is the use of CowSwap, which relies solely on on-chain liquidity and cannot leverage off-chain options like OTC desks or centralized exchanges, thus restricting potential solutions [14–16].

- **Efficiency and Service Quality Issues**: Many intent-based systems enforce strict on-chain processes, which are inherently slow and costly. This restriction of current solutions prevents the use of more efficient or cost-effective off-chain methods. For example, the automated market maker (AMM) model requires maintaining large amounts of liquidity on-chain, which ties up significant capital and incurs high costs. Similarly, cross-chain bridges require complex security mechanisms to ensure safe transfers, which increase both the time and financial cost of transactions. Additionally, due to the inability to handle asynchronous operations, service providers often have to bear higher service costs and potential risks. In a typical e-commerce scenario, a user might want to purchase an item, with the payment and delivery being asynchronous processes, reducing the risk for the service provider. However, current solutions require synchronous execution of intents, which increases the risk exposure for service providers. As a result, services are often slow and expensive, failing to meet the efficiency expectations of users accustomed to traditional financial systems [10, 18–20].

In summary, while intent-based systems represent a significant advancement in simplifying user interactions and enhancing efficiency within the blockchain ecosystem, they still require more flexible and comprehensive solutions. Addressing these limitations will pave the way for a more user-friendly and efficient blockchain experience, thereby driving mass adoption and unlocking the full potential of decentralized applications and services.

3.2 Optimistic Minimum Staking

The core of DAPPOS intent execution network is the Optimistic Minimum Staking (OMS) mechanism, an innovative solution perfectly tailored for intent-based tasks. OMS links each intent task to a predefined value and potential compensation in case of failure, allowing tasks to be executed before their outcomes are verified. In DAPPOS, intents are essentially framed as “provide C, want X, get Y when fail.” This means that if a task fails, the network employs a consensus process to penalize the responsible service provider, ensuring that either the user’s task is successfully executed or the user receives the predetermined compensation. Additionally, intents within the DAPPOS network require that the success of intent task execution can be clearly and objectively determined on-chain, eliminating subjective disagreements. This ensures transparency and consistency in the evaluation of task outcomes, further strengthening the reliability of the network. The critical innovation here is the asynchronous nature of C and X in this framework. The service provider can first receive C (the resources provided) from the user and then take the necessary steps to achieve X (the desired outcome), with the assurance that if they fail, Y (the compensation) will be provided to the user [24–27]. This asynchronous handling of C and X significantly expands the solution space for service providers. They are not constrained to synchronous, immediate actions but can leverage more complex, multi-step processes that might involve both on-chain and off-chain methods to fulfill the user's intent. The OMS mechanism offers significant advantages to both service providers and users. The key benefits of OMS include:

- **Generalizability of Intent Tasks**: OMS supports a wide variety of intent tasks, as long as they can be validated by on-chain rules. This broad applicability allows for numerous use cases, making the network versatile and adaptable. By leveraging on-chain validation, DAPPOS ensures that the execution of intents is transparent, secure, and consistent across different scenarios.
- **Expanded Solution Space**: OMS allows service providers to use any method to fulfill the user's intent as long as the final result is achieved. For instance, if a user wants to purchase a large amount of an asset, the service provider can choose the most efficient route, whether it's using on-chain exchanges, off-chain centralized exchanges like Binance, or over-the-counter (OTC) desks. This flexibility significantly broadens the range of possible solutions, enabling service providers to find the most cost-effective and timely methods [14–16].
- **Enhanced Capital Efficiency**: Service providers only need to stake a minimal amount of collateral relative to the total value of ongoing intent tasks, and they can take user's C(the resources provided) first and execute the intent later. This reduces the capital burden on service providers and allows them to utilize their funds more efficiently. They can freely deploy and utilize their working capital as long as they can fulfill the user's intent.
- **Guaranteed Execution Speed**: Execution speed is a critical component of the DAPPOS network, ensured through a predefined parameter within each intent. If a task is not completed within the specified time frame, it is automatically deemed a failure, and the user is compensated. This time-bound approach ensures that users are not left waiting indefinitely for their intents to be executed. It also incentivizes service providers to prioritize timely execution, knowing that delays will result in penalties and loss of collateral.
- **Scalability and Security**: The OMS system ensures that tasks are either completed successfully or the users are compensated adequately. This model guarantees that the network can handle a wide variety of tasks without compromising security. By focusing on the final outcome and using a robust consensus mechanism to validate task completion, DAPPOS maintains a high level of security while supporting scalable operations. The system's high tolerance for verification time and low-risk exposure ensure that tasks can be verified efficiently and securely, promoting a resilient and scalable network.

The OMS mechanism aligns the interests of service providers and users, fostering a reliable and efficient ecosystem. By ensuring that tasks are either completed successfully or compensated appropriately, OMS establishes a robust framework for intent execution, driving the adoption of intent-based interactions within the blockchain space.

3.3 Participants in Intent Execution Network

Before diving into the detailed workflow of the DAPPOS intent execution network, it is essential to introduce the three primary participants within the network:

Participants in Intent Execution Network

Figure 6: Participants in Intent Execution Network

Service Provider The service provider is the entity responsible for executing the user's intent. Service providers must stake collateral to take on tasks within the network, and the total value of the tasks they accept cannot exceed a certain proportion of their collateral. When a user's task is successfully completed and validated by the network, the reduction in the current total value of the service provider's tasks is acknowledged by other network participants, allowing the service provider to accept new tasks.

Execution Validator Execution validators play a crucial role in verifying whether a user's intent has been successfully executed. These validators must stake a significant amount of tokens as collateral within the network. By continuously participating in verification tasks, execution validators earn token rewards. If they fail to provide timely validation results or if they submit incorrect validation outcomes, they face penalties. Their role is vital to maintaining the integrity and accuracy of the network's operations.

Matcher Matchers are advanced execution validators with additional responsibilities. Besides verifying the execution results, matchers are also tasked with pairing and matching users with suitable service providers. This dual role requires matchers to stake even more assets within the network. Each service provider must select one or more matchers upon registration. Matchers ensure that the network operates smoothly by facilitating the interaction between users and service providers, enhancing the efficiency and reliability of the intent execution process [14–16]. These three roles—service providers, execution validators, and matchers—form the backbone of the DAPPOS intent execution network. Each participant plays a distinct and crucial part in ensuring that user intents are executed efficiently, accurately, and securely.

3.4 Workflow of Intent Execution Network

The DAPPOS intent execution network operates through a structured workflow, consisting of three main phases: Counterparty matching, execution, and verification.

Phase One: Counterparty Matching When a user defines their intent, they can query any matcher within the network. The matcher will return execution quotes from the connected service providers for the specified intent. Once the user selects

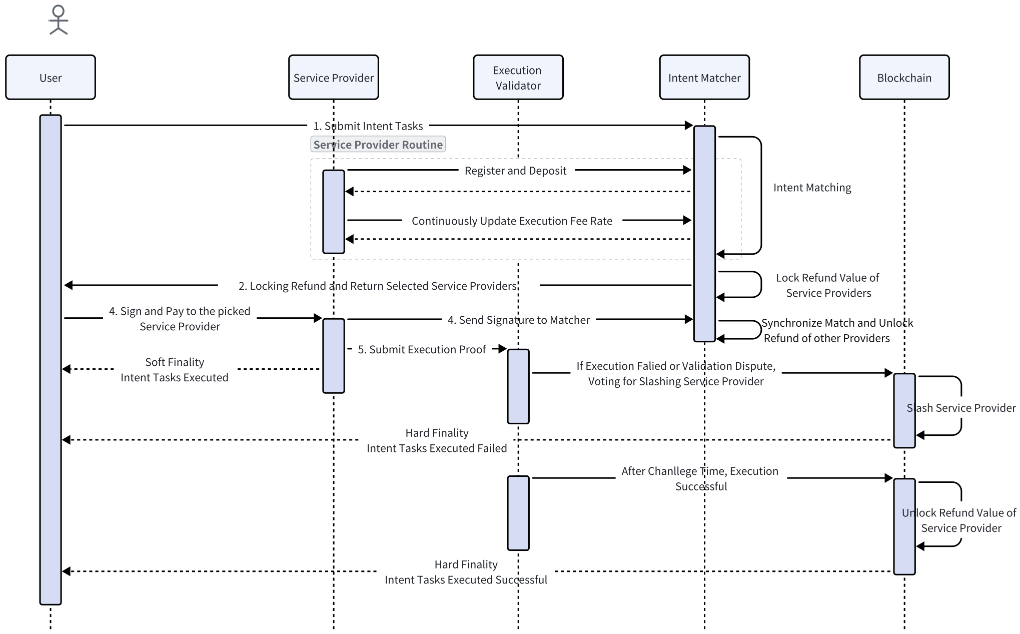

Figure 7: Workflow of Intent Execution Network

the most optimal quote, the intent deal is finalized. The user then only needs to monitor the success of their task; if it fails, they will receive compensation. To enhance user convenience, DAPPOS Web3 AI OS incorporates a suite of intent task frameworks that automatically formalize user intents and select the optimal service providers. As a result, users can enjoy seamless one-click execution without needing to delve into the intricacies of the Intent Execution Network.

Phase Two: Execution After the intent deal is finalized, the service provider will receive C(what user provides for the intent), and take responsibility for fulfilling the X(what user wants in the intent). DAPPOS Intent Execution Network does not care about how the service provider completes the intent task; it only cares about whether the result is successfully achieved. Unlike many other intent-related projects where intent solvers are limited to on-chain solutions, DAPPOS allows service providers to choose from a broader range of solutions. This flexibility may offer users better options.

Phase Three: Verification Service providers can proactively report the completion of their tasks. If they do not, execution validators will verify the success of the intent after the specified time within the intent. If the task is detected as failed, execution validators can issue a challenge to the network. All execution validators will then vote according to a Delegated Proof of Stake (DPOS) mechanism to reach a consensus. If the consensus determines that the intent task has failed, the service provider's collateral will be used to compensate the user and reward the challenger. If the consensus confirms that the intent task was successfully executed, the challenger will be penalized. If the task is successful, which means that nobody issues any challenge about the success, the corresponding task value for the service provider will be released.

Through these phases, DAPPOS ensures a seamless, efficient, and user-centric process for executing intents. The structured workflow not only simplifies user interactions but also maintains the network's integrity and reliability.

4 Innovations

DAPPOS represents a pioneering fusion of advanced AI orchestration and Web3-native execution, redefining how users interact with decentralized ecosystems. At its heart, the Intelligence Layer introduces groundbreaking innovations that deliver composable, adaptive intelligence far beyond the capabilities of generic AI models. These advancements, centered on the Multi-Agent Framework (MAF) and the Bubble Engine, enable precise, explainable reasoning tailored to Web3's volatility, from real-time market analysis to opportunity detection [2–9].

The MAF's stateful graph backbone, known as the StateGraph, serves as a foundational innovation, acting as the single source of truth for all system operations. By encapsulating variables such as messages, plans, artifacts, UI interrupts, outputs, errors, and flags, it ensures seamless state propagation across nodes, enabling deterministic routing and hierarchical control [2, 4, 5]. This structure powers the task_planner_node, which classifies user intents and directs them to specialized subgraphs—search for in-depth exploration, DeFi for strategy generation, and Opportunity for alpha identification—minimizing variability and fostering trust in high-stakes environments.

Within these subgraphs, innovations like the agentic search loop in the Search Subgraph exemplify autonomous collaboration: parallel agents (e.g., for web, social, and on-chain data) execute concurrently under a controller, with outputs merged, filtered, and refined through iterative enhanced review cycles until confidence thresholds are met. This results in transparent, cited summaries optimized for user interfaces [1–3, 7–9]. Similarly, tool-augmented generation in the DeFi and Opportunity Subgraphs integrates over 200 tools to simulate trades, validate risks, and produce executable plans, incorporating HCI interrupts for human oversight [2, 3]. Streaming mechanisms and safety rails further enhance interactivity, allowing real-time progress updates and pauses for clarification, creating a synergistic human-AI loop that outperforms siloed AI approaches [6].

Complementing MAF, the Bubble Engine's innovations in continuous reinforcement learning (RL) drive perpetual adaptation to Web3's fast-evolving landscape. Its Contextual Retrieval-Augmented Generation (RAG) employs a hybrid retriever that combines semantic and keyword search with domain-specific routing, deduplicating and reranking results to assemble provenance-first evidence packs. Citation-locked generation ensures traceability, while dynamic budgets scale retrieval based on uncertainty, blending episodic memory for recent interactions with long-term storage for verified knowledge [1,6–9]. The Compound Memory system layers durable Web3 fundamentals (e.g., protocol mechanics) with user-contributed "bubbles" from the task platform, accelerating evolution through real-time ingestion from sources like X and Binance Square. Robust misinformation detection—via on-chain cross-verification and community incentives—maintains integrity, turning potential deception into opportunities for collective refinement [10, 18–20].

In the Execution Layer, the Optimistic Minimum Staking (OMS) mechanism innovates intent fulfillment by enabling asynchronous operations, where service providers access user resources upfront while staking minimal collateral against failure penalties. This expands solution spaces beyond on-chain limits, incorporating off-chain efficiencies, and ensures guaranteed speed through time-bound execution. Supported by a Delegated Proof of Stake consensus among validators and matchers, OMS delivers institutional-grade reliability, hardened by millions of transactions.

Collectively, these innovations position DAPPOS as a transformative Web3 AI OS, unifying intelligence and execution to democratize complex crypto interactions and unlock ecosystem-wide value creation.

5 $DOS Tokenomics

5.1 Overview

DAPPOS is introducing $DOS, the native utility and governance token designed to support participation across the DAPPOS ecosystem. $DOS coordinates access, payments, staking, and governance within DAPPOS. $DOS is designed to align users, service providers, ecosystem contributors, and token holders around the long-term development of DAPPOS.

5.2 Utility

$DOS is the utility and governance token of DAPPOS. It is used to access selected premium services, pay transaction fees for intent executions, support staking-based service participation, and participate in protocol governance. Through these functions, $DOS connects the Intelligence Layer, the Execution Layer, AI service contributors, third-party service providers, and the broader DAPPOS community.

5.2.1 Subscription for Premium Features

$DOS enables users to subscribe to enhanced services within DAPPOS, including premium Bubble Engine capabilities. These capabilities may include proactive SOP(Standard Operating Procedures) optimization, customized AI workflow support, and selected advanced AI services for more specialized tasks. Users who stake $DOS may also become eligible for additional AI consumption credits, which can be used for selected premium AI services within DAPPOS according to product and protocol rules.

5.2.2 Payment of Transaction Fees

Users can pay transaction fees for intent executions using $DOS, benefiting from discounts compared to payments in other mainstream tokens, making it a cost-effective option within the ecosystem.

5.2.3 Staking Module

Ecosystem participants may be required to stake $DOS as performance collateral and service qualification. These participants include but are not limited to execution Layer service providers, execution validators, third-party AI agent providers, infrastructure partners, user support partners, and other service suppliers. The staking module supports network security, service reliability, and incentive alignment. Qualified participants may receive protocol-defined fees or rewards for services provided to the ecosystem. Token holders who do not provide services directly may delegate $DOS to eligible ecosystem participants under applicable protocol rules.

5.2.4 Protocol Governance

$DOS holders can participate in the ecosystem's future through on-chain governance, such as proposing and voting on protocol decisions that shape DAPPOS's evolution according to community vision and needs.

5.3 Allocation and Vesting Schedule

The total supply of token $DOS is 1,000,000,000. DAPPOS is committed to transparency regarding the distribution and release of $DOS. The allocation and vesting schedule is designed to ensure long-term alignment between the team, investors, ecosystem contributors, service participants, and the broader community.

Category	**Allocation**	**Vesting Schedule**
Team	20%	12-month cliff, followed by linear distribution over the subsequent 48 months.
Investors	22.5%	12-month cliff, followed by linear distribution over the subsequent 48 months.
Ecosystem	20%	5% initial release at TGE, followed by 3.5% monthly release for 12 months, with the rest linear over 48 months.
Treasury	20%	9% initial release at TGE, followed by 2% monthly release for 12 months, with the rest linear over 48 months.
Marketing	11.5%	3% initial release at TGE, remainder on a 12-month cliff followed by linear vesting over 48 months.
Airdrop	6%	3% initial release at TGE, followed by 3% monthly release over the next 3 months.

Table 1: $DOS Allocation and Vesting Schedule

5.3.1 Ecosystem Bootstrapping

Initial unlocks in the first 12 months after TGE are entirely ecosystem-based and community-based. 20% of the total token supply unlocks at TGE. The $DOS token airdrop accounts for 6% of the total supply. Of this, 3% will be distributed at TGE, with the remaining 3% distributed monthly from TGE+1 through TGE+3. Primary recipients of the airdrop fall into the following categories: participants in DAPPOS Execution and Intelligence Layer testing, recognized contributors and KOLs within the DAPPOS ecosystem, subscribers of DAPPOS AI, and holders of DAPPOS Bubble NFT

5.3.2 Core Contributors and Investors

None of the team allocation or the investor allocations will be unlocked at TGE. The team and investor allocation will unlock gradually over 48 months following a 12-month cliff. This structure is designed to align core contributors and early supporters with the long-term development, maintenance, and expansion of the DAPPOS ecosystem.

5.4 Disclaimer

$DOS does not represent ownership, income rights, or any form of security interest. It does not guarantee future price appreciation or market performance. Its sole purpose is to enable utility and participation within the DAPPOS ecosystem.

6 Roadmap

DAPPOS's roadmap follows the architecture described throughout this whitepaper: the Intelligence Layer, the Execution Layer, and a growing ecosystem of AI-powered services built around user intents. DAPPOS starts from Web3, where users face complex information, fragmented tools, and high execution barriers. However, the same architecture is not limited to Web3. A system that researches, plans, optimizes SOPs, and coordinates execution can extend into broader domains where users need domain-specific intelligence and reliable task fulfillment.

6.1 Generalized Intelligence

DAPPOS will continue extending the Intelligence Layer from crypto-native tasks toward broader research and decision-support scenarios. This includes domains such as U.S. equities, market intelligence, business research, and other information-rich fields where users need timely data, domain-specific reasoning, and structured outputs. The goal is to make DAPPOS useful not only for Web3-native users, but also for broader users who need AI to research, analyze, and organize complex decisions.

6.2 Premium Intelligence and SOP Optimization

DAPPOS will continue improving the Bubble Engine's ability to build, test, and optimize task-specific SOPs. This stage focuses on premium intelligence services, including proactive SOP optimization, customized workflow tuning, and more stable AI outputs for specialized tasks. As more SOPs are created and improved, DAPPOS can reduce the burden on users to operate AI tools directly.

6.3 AI-Assisted Application and Business Building

DAPPOS will further support users and developers in turning goals into usable software and business workflows. This includes AI-assisted creation of Web3 applications, websites, dashboards, lightweight AI agent applications, and other task-specific interfaces. For users who want to build Web3-related products or workflows, DAPPOS can help translate business intent into application logic, interface structure, execution flow, and service configuration. The broader direction is to move from AI-assisted research and planning toward AI-assisted building, so that more users can create applications and workflows without needing to manage every technical layer themselves.

6.4 Ecosystem Service Network

As DAPPOS expands into more domains, the ecosystem will require a broader service network. This network may include Execution Layer service providers, execution validators, third-party AI agent providers, infrastructure partners, workflow specialists, and user support partners. These participants help DAPPOS support more use cases while keeping the protocol modular and scalable. $DOS staking will serve as a qualification and incentive-alignment mechanism for eligible ecosystem participants, supporting more reliable service delivery across the network.

6.5 Protocol Governance

As the ecosystem matures, DAPPOS will progressively expand governance participation for $DOS holders. Governance may cover protocol parameters, staking rules, service-provider participation, ecosystem standards, and future extensions of the DAPPOS AI operating layer. In the long term, DAPPOS aims to become a generalizable AI operating system that starts from Web3, extends into broader research and application-building scenarios, and enables users to move from intent to intelligence, and from intelligence to execution.

References

[1] Patrick Lewis et al. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. NeurIPS, 2020.

[2] Shunyu Yao et al. ReAct: Synergizing Reasoning and Acting in Language Models. arXiv:2210.03629, 2022/2023.

[3] Timo Schick et al. Toolformer: Language Models Can Teach Themselves to Use Tools. arXiv:2302.04761, 2023.

[4] Qianyang Wu et al. AutoGen: Enabling Next-Gen LLM Applications via Multi-Agent Collaboration. arXiv:2308.08155, 2023.

[5] Guohao Li et al. CAMEL: Communicative Agents for “Mind” Exploration of LLM Society. ICLR, 2023.

[6] Charles Packer et al. MemGPT: Towards LLMs as Operating Systems. arXiv:2310.08560, rev. 2024.

[7] Weizhi Fan et al. A Survey on RAG Meeting LLMs. ACM, 2024.

[8] Y. Gao et al. Retrieval-Augmented Generation for LLMs: A Survey. arXiv:2312.10997, 2023.

[9] Y. Huang & J. Huang. A Survey on Retrieval-Augmented Text Generation for LLMs. arXiv:2404.10981, 2024.

[10] Hayden Adams et al. Uniswap v3 Core Whitepaper. 2021.

[11] Anoma Foundation. An Introduction to Intents and Intent-Centric Architectures. 2023.

[12] Essential. Essential: The First Declarative Blockchain. 2024.

[13] The Block. "Intent-centric Ethereum L2 Essential raises $11M." Aug 2024.

[14] CoW Protocol Docs. Batch Auctions and Trade Intents. 2024–2025.

[15] CoW Protocol Docs. Solvers and Batch Auction Mechanism. 2024–2025.

[16] 1inch. Fusion / Fusion+ (intent-based swaps). Docs, 2024–2025.

[17] ERC-4337. Account Abstraction. Official resources, 2023–2025.

[18] J. Notland et al. SoK: Cross-Chain Bridging Architectural Design Flaws and Mitigations. arXiv:2403.00405, 2024.

[19] ACM Digital Library. Security of Cross-Chain Bridges: Attack Surfaces and Defenses. 2024.

[20] R. Fritsch et al. Concentrated Liquidity in Automated Market Makers. arXiv:2110.01368, 2021.

[21] Alchemy. What is Account Abstraction (ERC-4337)? 2023.

[22] Stackup. What is EIP-4337? 2025.

[23] AfterDark Labs. EIPs Explained: EIP-4337. 2023.

[24] Daniel Larimer. Delegated Proof of Stake (DPoS) Whitepaper. 2014.

[25] BitShares Docs. Delegated Proof of Stake (DPOS). 2014–2025.

[26] W. Li et al. Delegated Proof of Stake Consensus Based on Node Trust. Appl. Sci., 2023.

[27] Q. Hu et al. An Improved DPoS Consensus Algorithm. Procedia Comp. Sci., 2021.