

White Paper: Unibase: Decentralized AI Memory Layer for Autonomous Agents

Unibase Community, www.unibase.com

September, 2025

Abstract

Unibase is the first high-performance decentralized AI memory layer designed to empower autonomous AI agents with persistent memory, cross-platform interoperability, and verifiable data sovereignty. As AI systems increasingly move toward agent-based models, the lack of standardized memory infrastructure, transparent data ownership, and seamless agent collaboration has become a critical bottleneck limiting the potential of decentralized AI. Unibase addresses these fundamental challenges through three integrated core modules: **Membase** (decentralized memory layer), **AIP Protocol** (Agent Interoperability Protocol), and **UnibaseDA** (high-performance data availability layer). Together, these components paves the way for a new era of open, secure, and scalable AI ecosystems for the Open Agent Internet — a composable, trustless network where intelligent agents can learn, collaborate, and evolve autonomously across platforms while maintaining data sovereignty and verifiable interactions.

1 Introduction

1.1 The AI Agent Evolution

The rise of large language models (LLMs) and autonomous AI systems in recent years has unlocked unprecedented capabilities in machine learning and natural language processing. These models have enabled complex forms of reasoning and language understanding, propelling AI beyond static, single-purpose tasks toward interactive systems capable of contextual conversation, decision-making, and task execution. As these AI agents have grown more capable, the way users interact with them has also changed, shifting from one-off, question-and-answer sessions to multi-turn dialogs and extended tasks that span multiple sessions or platforms. This transformation represents a crucial step in the evolution of AI, marking a transition from short-lived, stateless applications to long-lived, context-aware autonomous agents capable of adapting over time.

However, effectively deploying these advanced AI agents presents three major architectural hurdles. First, the prevalent “stateless” design in today’s AI deployments means that agents reset at every user interaction, failing to preserve context or learn from historical experience. Second, these agents are largely isolated within proprietary platforms, limiting the sharing of data, insights, or capabilities with other AI agents and tools. Third, control over user data, agent behavior, and governance aspects is often centralized within key service providers—limiting user choice, introducing single points of failure, and constraining the innovation potential of a broader community. Addressing these issues is essential to unlock the full promise of “truly autonomous” AI systems and to lay the groundwork for an open, decentralized agent network.




Figure 1: Connecting multiple entities through Unibase.

In pursuit of such a future, Unibase propose the concept of the Open Agent Internet (OAI) shown in Figure 1. OAI provides modern AI agent infrastructure for entities around the globe to connect for communications. At its core, the OAI envisions a network of intelligent entities that maintain persistent memories and can communicate and collaborate fluidly across institutional and infrastructural boundaries. By leveraging principles from Web3 and distributed systems, the OAI strives to ensure robust data sovereignty, user ownership of AI-driven interactions, and cryptographically verifiable decisions. Figure 1 illustrates an example of how multiple entities might connect through a decentralized memory layer, called Unibase, to enable interoperability and long-term information persistence. This decentralized approach not only enhances system resilience—by avoiding a single point of failure—but also empowers users and developers to build new kinds of agent collaboration frameworks under open governance models.

1.2 Limitations of Current AI Agent Systems

Although many contemporary AI systems exhibit advanced capabilities, they are constrained by several limiting factors that hamper their ability to function autonomously and collaboratively at scale. First, the absence of persistent memory across interactions constrains agents to a stateless operational model. Each conversation or task begins with no recollection of previous episodes, impeding the agent’s ability to learn incrementally or accumulate domain knowledge over extended periods. As interactions grow in complexity, reprocessing entire conversation histories becomes computationally expensive, raising both infrastructure costs and latencies that degrade responsiveness.

Second, existing AI platforms remain largely isolated from one another. Proprietary APIs, mismatched data structures, and heterogeneous communication protocols hinder agents from discovering and utilizing external services, tools, or even other complementary AI models. Current protocols that do enable inter-agent messaging—such as Google’s Agent-to-Agent (A2A) protocol [29]—do not incorporate facilities for persistent memory, secure identity management, or robust reputation tracking. Likewise, Anthropic’s Model Context Protocol (MCP) [12] focuses on integrating LLMs with tools but does not address the deeper need for ongoing inter-agent connectivity or cryptographically verifiable data storage. These platform silos limit opportunities for synergistic knowledge-sharing and reduce overall innovation velocity by forcing developers to constantly re-implement integration layers.

Finally, data governance remains concentrated in corporate or platform-centric silos, undermining user sovereignty. In currently dominant models, users have limited visibility into how data is utilized, and generally have no control over how it is shared between third-party services or used to retrain proprietary models. This not only raises privacy concerns but also creates single points of failure that threaten data availability and system resilience. Moreover, attempts to leverage blockchain or Web3 solutions for decentralized data storage face inherent performance bottlenecks, such as high latency and limited throughput, rendering them unsuitable for real-time AI inference and large-scale ML workloads. Without a high-performance, privacy-preserving, and decentralized infrastructure layer, these fundamental challenges in memory persistence, interoperability, and data governance remain unaddressed.

1.3 Unibase: Decentralized AI Memory Layer for Autonomous Agents

To overcome these challenges, the report presents Unibase, a decentralized memory layer designed to serve as the backbone for next-generation AI agents. Unibase aims to couple the strengths of Web3 decentralization—such as tamper-proof data storage and cryptographic identity verification—with the performance demands of modern ML systems. Specifically, it provides cryptographically secured, low-latency, and high-throughput data operations to ensure that AI agents can access, update, and share knowledge in real time. This decentralized memory layer supports continuous learning by persisting contextual information, chat histories, task progress, and even model updates across sessions and platforms.

Crucially, Unibase offers standardized interoperability through universal agent-to-agent (A2A) and agent-to-tool (A2T) protocols, enabling a wide array of AI services to seamlessly communicate. Agents from different ecosystems can share structured memory records, exchange tasks, and chain capabilities without ad-hoc integration layers. Given the decentralized nature of Unibase, users retain control over how data is stored, shared, or monetized—preserving data sovereignty while mitigating centralized points of failure. In addition, cryptographic proofs of data integrity allow for transparent verification of agent actions and states, providing a mechanism to enhance trust between autonomous agents, service providers, and end-users.

Internally, Unibase employs a scalable, modular architecture that accommodates exabyte-scale data while meeting latency requirements of under 100 ms for real-time inference workloads. By leveraging distributed consensus protocols optimized for high concurrency, Unibase can seamlessly integrate with existing HPC clusters and data centers, bridging traditional cloud infrastructure with on-chain proofs of data authenticity. With this foundation, developers gain an ecosystem that not only accommodates today’s machine learning demands but also anticipates future advances in decentralized AI, where thousands—or millions—of agents need to securely coordinate tasks, exchange learning insights, and preserve knowledge across platforms.

In essence, Unibase positions itself as the missing infrastructure piece necessary for achieving the broader vision of the Open Agent Internet. By addressing memory statelessness, platform isolation, and centralization in current AI solutions, Unibase paves the way for building open, composable, and verifiably secure agent-based applications at global scale. In addition, Unibase natively integrates the emerging ERC-8004 “Trustless Agents” standard, exposing Ethereum-compatible identity, reputation, and validation registries so that agents built on Unibase can participate in open, cross-organizational agent economies on any EVM-compatible chain.

1.4 Report Organization

This report is organized as follows:

- **Section 2** provides an in-depth overview of the Unibase architecture, detailing its modular design and the functionalities of its core components: Membase, AIP Protocol, and UnibaseDA.
- **Section 3** delves into the workings of Membase, explaining its role in decentralized identity and memory management, and how it ensures persistent, secure, and verifiable agent memory across platforms.
- **Section 4** explores the AIP Protocol, highlighting its capabilities in enabling seamless agent interoperability and communication, as well as its compatibility with existing protocols and integration with Web3 technologies.
- **Section 5** focuses on UnibaseDA, the decentralized storage and data availability layer, describing its novel architecture that supports high throughput and scalability while maintaining stringent security and verification standards.
- **Section 6** presents the phased deployment strategy for Unibase, outlining the timeline from mainnet launch through the implementation of cross-platform memory synchronization, including key milestones for governance mechanisms, zero-knowledge verification systems, and the release of AIP 2.0 for full agent interoperability.
- **Section 7** discusses the Unibase ecosystem, showcasing practical applications and use cases, as well as partnerships and integrations with other agent frameworks and platforms.
- **Section 8** examines the UB token economics and incentive mechanisms, detailing how protocol fees, vote-escrowed governance (veUB), memory node staking, and knowledge mining rewards create a sustainable economic model that aligns stakeholder interests while enabling horizontal scalability to millions of nodes and fostering long-term ecosystem growth.
- **Section 9** concludes the report by summarizing the contributions of Unibase to the field of decentralized AI and outlining future directions for research and development.

2 Unibase Architecture

2.1 Overview

As depicted in Figure 2, Unibase employs a layered and modular design that separates key system functionalities while maintaining cohesive performance at scale. This architecture is organized into three primary layers: (1) the Application Layer, represented by Membase; (2) the Service Layer, governed by the AIP Protocol; and (3) the Infrastructure Layer, realized by UnibaseDA. Such a hierarchical structure allows each layer to be optimized independently, yet still achieve the end-to-end throughput, security, and scalability required by autonomous AI agents.




Figure 2: Unibase architecture.

At a high level, agents make use of the AIP Protocol for secure and standardized communication, rely on Membase for persistent identity and memory management, and draw on the high-performance storage capabilities of UnibaseDA. Meanwhile, the underlying blockchain component ensures robust identities, governance mechanisms, and on-chain verifiability—without becoming a bottleneck for real-time data access. By distributing these responsibilities across separate but tightly integrated subsystems, Unibase provides an extensible platform that promotes interoperability, transparency, and decentralized control.

2.2 Membase: Decentralized Identity and Memory Management

Membase functions as the foundational layer for identity assignment and persistent memory storage within Unibase. It provides each AI agent with a unique, verifiable identifier via its Identity Registration subsystem, enabling secure recognition and interaction across decentralized environments. Beyond registration, Membase also handles Configuration & History Tracking, preserving agent configurations, activity logs, and interaction histories to deliver seamless context continuity across sessions or platforms.

In addition, Membase implements robust Authorization Mechanisms to safeguard data access and interaction privileges. These permission structures ensure that agents are granted only the level of access they require, supporting nuanced, context-sensitive policy controls. By consolidating identity management, historical records, and permission enforcement, Membase serves as the anchor point for consistent and verifiable agent experiences across the Unibase ecosystem.

2.3 AIP Protocol: Agent Interoperability and Communication

Situated above Membase, the AIP Protocol defines standardized communication, authentication, and authorization procedures for AI agents within Unibase. Designed for Web3 compatibility, the AIP Protocol supports both MCP- and gRPC-based interactions, enabling real-time agent-to-agent messaging and shared memory references. Its standardized workflows describe universal message formats and data structures, reducing friction in cross-platform or cross-vendor scenarios.

Crucially, AIP secures agent interactions through decentralized authorization, where cryptographic signatures and flexible access-control checks replace monolithic central servers. Multi-layered trust builds on stored on-chain credentials, zero-knowledge proofs, and robust permission schemes to further protect agent-to-agent connections. Through these trust primitives, AIP estab-

lishes a privacy-preserving and tamper-resistant framework that bolsters confidence in distributed agent interactions across the Unibase network.

2.4 UnibaseDA: Decentralized Storage & Data Availability

At the infrastructure level, UnibaseDA is responsible for providing high-throughput, verifiable storage and retrieval of AI data. To maintain censorship-resistance and durability, UnibaseDA uses distributed storage techniques such as erasure coding: data is intelligently fragmented, distributed, and safeguarded across numerous decentralized nodes. This approach ensures consistent data availability while minimizing risks associated with single points of failure. On-chain verifiability underpins data integrity checks and transparency by linking data availability proofs to blockchain consensus and zero-knowledge verification. Efficient retrieval builds on this foundation by employing decentralized indexing and streamlined lookups, keeping read latencies low even at large scales. By bringing together decentralized data storage with robust on-chain proofs, UnibaseDA ensures reliable, high-performance operations for agents demanding sub-100 ms response times and throughput exceeding 100 GB/s, effectively eliminating the need for fragile, centralized storage silos.

3 Membase: Decentralized Identity and Memory Management

AI agents have demonstrated remarkable capabilities in complex reasoning, natural language processing, and task execution. However, current agent architectures face fundamental limitations when deployed in distributed, multi-agent environments. The primary challenge lies in the ephemeral nature of agent memory, where conversation history, learned knowledge, and operational context are lost between sessions or platform migrations.

Traditional centralized approaches to agent memory storage introduce single points of failure, privacy concerns, and vendor lock-in problems that are incompatible with the principles of decentralized autonomous systems. Furthermore, the lack of standardized identity verification mechanisms prevents secure inter-agent communication and coordination, limiting the potential for collaborative multi-agent workflows.

The emergence of Web3 technologies and blockchain infrastructure provides an opportunity to address these limitations through decentralized storage, cryptographic identity verification, and transparent coordination mechanisms. However, existing blockchain solutions are not optimized for the high-throughput, low-latency requirements of AI agent memory systems, necessitating a specialized architecture that bridges the gap between traditional AI infrastructure and decentralized technologies.

Unibase introduces Membase, a comprehensive decentralized AI memory layer that addresses these challenges through three primary innovations: (1) a dual-layer identity system combining on-chain verification with efficient off-chain operations, (2) a hierarchical memory architecture supporting multiple storage modalities from real-time conversations to long-term knowledge bases, and (3) blockchain-integrated task coordination enabling verifiable multi-agent collaboration.

3.1 Membase Design

3.1.1 Design Principles

Membase is architected around three fundamental design principles that address the unique requirements of decentralized AI agent systems:

Verifiable Identity: All agent operations must be cryptographically verifiable and traceable to prevent impersonation and ensure accountability in multi-agent interactions. This principle drives the dual-layer identity architecture that separates ownership control from operational identity.

Persistent Memory: Agent memory must survive platform migrations, service interruptions, and extended offline periods while maintaining data integrity and accessibility. This requirement necessitates the hierarchical storage approach with multiple redundancy layers.

Scalable Performance: The system must support high-throughput agent operations without compromising decentralization or security, requiring careful optimization of data access patterns and caching strategies.




Figure 3: Membase architecture.

3.1.2 Membase Overview

As shown in Figure 3, the Membase architecture consists of three primary layers that work in concert to provide comprehensive agent memory and coordination services:

3.2 Agent Hub

The Agent Hub serves as the central coordination point for all agent operations, managing the full lifecycle of agent interactions through three specialized sub-modules: (1) Link Hub facilitates remote interaction services, enabling agents to establish secure connections with external services, other agents, and blockchain networks. The Link Hub implements protocol translation layers that allow agents to communicate across different network protocols while maintaining identity verification throughout the interaction chain. (2) Config Hub manages agent configurations including identity credentials, operational parameters, permission sets, and service discovery information. Config Hub maintains a distributed registry that enables efficient agent discovery while supporting dynamic configuration updates without service interruption. (3) Memory Hub provides the primary interface for persistent data operations, handling conversation storage, knowledge base management, and interaction history tracking. The Memory Hub implements sophisticated caching strategies to optimize data access patterns while ensuring data consistency across distributed storage layers.

3.3 Membase Blockchain Contract

The blockchain layer provides cryptographic foundations for the entire system through on-chain storage of critical identity and permission data. Unlike traditional blockchain applications that attempt to store all operational data on-chain, Membase employs a hybrid approach where only essential verification data resides on-chain while operational data is stored in the decentralized availability layer.

The chain layer maintains agent identity registrations, permission mappings, and task coordination smart contracts. This design ensures that all agent operations can be cryptographically verified while avoiding the performance and cost limitations of storing large datasets directly on blockchain networks.

3.3.1 Hierarchical Memory Layers

Membase identity system addresses the fundamental challenge of secure agent identification through a novel dual-layer architecture that separates ownership control from operational identity. This design enables flexible agent management while maintaining strong security guarantees.




Figure 4: Membase workflow.

Agent Owner Layer: The agent owner layer represents the ownership control layer, functioning as the root identity for one or multiple agent instances. Each account is cryptographically secured and serves the following critical functions:

- **Ownership Control:** The account establishes clear ownership relationships for all associated agent identities, enabling secure delegation of permissions and centralized management of multiple agents under a single ownership umbrella.
- **On-Chain Registration:** Account registration occurs directly on the blockchain, providing immutable proof of identity establishment and ownership claims. The registration process includes cryptographic key generation, identity commitment, and initial permission set definition.
- **Off-Chain Verification:** The account credentials enable signing and verification of off-chain operations, allowing agents to perform high-frequency operations without incurring blockchain transaction costs while maintaining cryptographic verifiability.
- **Data Management:** The account layer aggregates data from all associated agent identities, enabling seamless data migration, backup operations, and cross-agent data sharing within the same ownership domain.

Agent Instance Layer: The agent instance layer serves as the operational identity layer, providing globally unique identifiers for individual agent instances. This layer is optimized for high-frequency interactions while maintaining traceability to the ownership layer:

- **Unique Identification:** Each agent instance is globally unique and verifiable on-chain, preventing identity collision and enabling reliable agent addressing across distributed networks.
- **Primary Interaction Identity:** The `MEMBASE_ID` serves as the primary identifier for all inter-agent communications, service requests, and task participation, providing a stable identity that persists across platform migrations.
- **Scalable Operations:** Multiple agent instances can operate under a single agent owner layer, enabling horizontal scaling of agent operations while maintaining centralized ownership control.

3.4 Memory Management

Membase implements a sophisticated memory management architecture that addresses the diverse storage requirements of AI agents through multiple specialized memory modalities. This approach recognizes that different types of agent data have distinct access patterns, persistence requirements, and performance characteristics.

3.4.1 Conversation Memory

The conversation memory subsystem manages real-time interaction data through the MultiMemory and BufferedMemory components, each optimized for different usage patterns:

MultiMemory Management: The MultiMemory system enables agents to maintain multiple concurrent conversation threads while providing efficient access to historical interactions. Each conversation is identified by a unique `conversation_id`, enabling agents to context-switch between different interaction contexts without data loss. The system implements intelligent memory organization that groups related messages and maintains conversation metadata including participant identities, timestamps, and interaction classifications. This structure enables sophisticated queries such as retrieving all conversations with specific agents or finding conversations containing particular topics.

BufferedMemory Operations: The BufferedMemory system provides a simplified interface for agents that primarily operate in single-conversation modes. This system optimizes for sequential access patterns and implements efficient buffering strategies that balance memory usage with access performance. Both memory systems support automatic synchronization with the decentralized storage layer, ensuring that conversation data persists beyond agent session boundaries while maintaining real-time access performance for active conversations.

3.4.2 Knowledge Base Integration

The knowledge base subsystem addresses the challenge of managing structured information that agents need for reasoning and decision-making. The Chroma KnowledgeBase implementation provides a sophisticated vector storage system optimized for semantic search and retrieval:

Document Management: The system supports ingestion of structured documents with rich metadata annotation, enabling agents to maintain comprehensive knowledge repositories that support both factual lookup and semantic reasoning operations.

Vector Embeddings: Documents are automatically converted to vector embeddings using state-of-the-art embedding models, enabling semantic similarity search that goes beyond keyword matching. This capability allows agents to retrieve relevant information even when queries don't exactly match stored content.

Distributed Storage: Knowledge bases automatically synchronize with the decentralized storage layer, ensuring that agent knowledge persists across platform migrations while supporting collaborative knowledge sharing between agents operating under the same account.

3.5 Conclusion

This chapter presents Membase, a comprehensive decentralized AI memory layer that addresses critical infrastructure needs for autonomous agent systems. Through the combination of dual-layer identity architecture, hierarchical memory management, and blockchain-integrated task coordination, Membase provides the foundational infrastructure necessary for the emergence of truly autonomous AI agents operating in decentralized environments.

4 AIP: Agent Interoperability and Communication

The rapidly expanding field of AI agent development is currently experiencing significant fragmentation. Existing protocols, such as Anthropic's MCP and Google's A2A protocol, have critical limitations, notably a lack of cross-platform identity, decentralized memory, and verifiable authorization. These shortcomings severely impede genuine multi-agent collaboration and hinder the realization of a truly open agent internet.

Existing agent protocols, while serving specific purposes, fall short in addressing these critical interoperability gaps: First, MCP (Anthropic): Primarily designed for integrating LLMs with tools and data, MCP lacks inherent support for cross-agent communication, built-in memory capabilities, verifiable on-chain identity, and decentralization. This confines agents to more siloed operations, limiting their ability to engage in complex, multi-party interactions. Second, A2A (Google): While focused on agent-to-agent communication, A2A similarly lacks crucial features such as seamless tool integration, persistent memory support, on-chain identity, and decentralized architecture. Its design does not inherently support the verifiable, trustless interactions increasingly demanded by advanced AI applications.

The fundamental gaps in both traditional protocols—specifically the absence of cross-platform identity, decentralized memory, and verifiable authorization—lead to fragmented, insecure, and

non-persistent agent interactions. This concentration of AI power and infrastructure in the hands of a few entities, as observed in the broader AI landscape, can stifle competition, impede progress, and restrict who can develop and benefit from cutting-edge AI applications. The current state describes a closed or fragmented agent ecosystem, where trust is often centralized and data ownership is ambiguous.

The following table visually demonstrates AIP’s comprehensive advantages over existing protocols, highlighting its unique value proposition and the breadth of its capabilities.

Feature	MCP (Anthropic)	A2A (Google)	AIP (Unibase)
Primary Focus	LLM & tool/data integration	Agent-to-agent communication	Full agent interoperability + tool access
Cross-Agent Communication	✗	✓	✓
Tool Integration	✓	✗	✓
Built-in Memory Support	✗	✗	✓ (via Membase)
On-Chain Identity & Auth	✗	✗	✓ (via ZK + blockchain)
Agent/Tool Discovery	✓	✗	✓ (built-in discovery & registry)
Protocol Compatibility	MCP only	A2A only	✓ (MCP + gRPC compatible)
Decentralization	✗	✗	✓ (Web3-native)

Table 1: Protocol Comparison: AIP vs. MCP vs. A2A

In response to these pervasive challenges, Unibase has introduced the Agent Interoperability Protocol (AIP), heralded as the first Web3-native multi-agent communication standard. AIP is designed to overcome the critical limitations of existing protocols by providing a comprehensive, full-stack solution for agent interoperability. The core promise of AIP is to enable agents to establish verifiable on-chain identities, share decentralized memory, and collaborate seamlessly across diverse platforms. Its Web3-native design is not merely a buzzword; it represents a fundamental architectural choice that leverages the core principles of decentralized technologies. This foundation ensures transparency, enhanced security, and true data ownership for agents and their interactions. By adopting a Web3-native approach, AIP directly addresses the trust deficits, data integrity issues, and censorship vulnerabilities that are inherent in centralized systems. This positions AIP as a critical infrastructure for the future of AI where trust, transparency, and verifiable interactions are paramount.

4.1 AIP Design

As shown in Figure 5, this section provides a detailed breakdown of AIP’s underlying technology and architecture, explaining how its core components work together to deliver the promised functionalities.

4.1.1 Design Principles

The design and development of AIP are guided by several core principles that underpin its robust and future-proof architecture:

- **Decentralization:** Built fundamentally on Web3 principles, AIP ensures that control, data, and identity are distributed across a network, thereby eliminating single points of failure, mitigating censorship risks, and promoting a more resilient ecosystem.
- **Interoperability:** AIP is meticulously designed for compatibility with both existing and emerging protocols, including MCP and gRPC. This forward-thinking approach ensures broad adoption and seamless integration into the current diverse AI landscape, allowing agents built on other protocols to readily leverage AIP’s advanced features.
- **Security & Trustlessness:** By leveraging blockchain technology for identity management and Zero-Knowledge Proofs (ZK) for authorization, AIP prioritizes verifiable security and




Figure 5: Demonstration of AIP Architecture.

enables trustless interactions among agents. This cryptographic assurance is vital for sensitive applications.

- **Persistence & Scalability:** Through its integration with Membase and UnibaseDA, AIP provides robust, scalable, and persistent memory solutions for agents, ensuring their long-term knowledge and operational context are preserved.
- **Autonomy:** By enabling direct LLM integration and comprehensive tool access, AIP empowers agents with enhanced autonomy, allowing them to execute complex workflows and make informed decisions without constant human oversight.

4.1.2 AIP Overview

The proposed system consists of three core layers: the AIP Runtime Layer, which serves as the central orchestration component for cross-agent and tool interactions, the Identity layer for secure agent registration and authentication, and the Memory layer for persistent and verifiable conversation storage. As illustrated in the figure, the LLM Agent Client communicates with the AIP Runtime Layer through gRPC and SSE/IO protocols. The runtime layer manages bidirectional communication with external Agents via gRPC and integrates with MCP tools through SSE/IO connections. Unlike conventional stateless agent architectures that discard conversation context after each interaction, Unibase AIP maintains persistent state through three specialized subsystems: the identity module handles secure agent onboarding and authentication, the Memory module provides persistent and verifiable storage of conversation history and prompts, and the DA layer ensures reliable storage with zero-knowledge proofs for privacy preservation. This architecture enables seamless multi-agent collaboration while maintaining conversation continuity across sessions, allowing agents to leverage historical context and maintain long-term memory without compromising security or verifiability.

4.2 AIP Protocol

The blockchain task coordination system enables transparent, verifiable collaboration between agents without requiring centralized coordination authorities. This system addresses the fundamental challenge of establishing trust and ensuring fair compensation in multi-agent collaborative workflows.

4.2.1 AIP Contract

The task coordination system is built on smart contracts that implement standardized protocols for task creation, assignment, execution verification, and reward distribution:




Figure 6: Agent task execution in AIP.

Task Registry: Tasks are registered on-chain with comprehensive metadata including requirements, success criteria, reward amounts, and deadline specifications. The registry provides a transparent marketplace where agents can discover available tasks that match their capabilities.

Agent Registration: Agents must register their capabilities and stake collateral to participate in the task system, creating economic incentives for honest behavior and providing mechanisms for dispute resolution.

Execution Verification: Task completion is verified through multiple mechanisms including cryptographic proofs, peer review, and automated validation where applicable. This multi-layered approach ensures that rewards are distributed fairly while minimizing the potential for gaming the system.

4.2.2 Coordination Protocols

The system implements standardized protocols that enable complex multi-agent workflows while maintaining transparency and verifiability:

Task Decomposition: Complex tasks can be automatically decomposed into subtasks that can be executed by different agents, enabling parallel execution and specialization while maintaining overall coordination.

Progress Tracking: All task execution steps are recorded on-chain, providing transparency for task creators and enabling sophisticated analytics about agent performance and system efficiency.

Dispute Resolution: Automated dispute resolution mechanisms handle conflicts between task creators and executing agents, using cryptographic evidence and peer review to make fair determinations.

4.3 AIP Task Execution

Figure 6 illustrates a multi-agent task execution framework consisting of four main components: task, agent, agent hub, and chain. The system operates through a structured sequence of interactions across four distinct phases. During the register phase, both tasks and agents register themselves with the agent hub, which subsequently synchronizes this registration information with the chain component. The agent hub also performs configuration updates to maintain system state consistency. The task initiation phase begins when a task registers itself and provides task




Figure 7: Agent-to-agent communication in AIP.

information to the agent hub. The system then facilitates the agent’s participation through a ”join task onchain” mechanism, where agents submit requests to join specific tasks. Upon approval, permission is granted, enabling the transition to the execution phase. In the task execution phase, the agent hub initiates the task job and coordinates the actual work execution. Agents perform their assigned tasks and submit results back to the agent hub, which serves as the central coordination point for task management and result aggregation. Finally, during the task termination phase, the system evaluates the submitted results and processes the final outcomes. Upon successful completion, the agent hub submits the final result to the chain component, which then distributes task rewards to participating agents, completing the full task lifecycle. Throughout this process, the agent hub functions as the central orchestrator, managing agent registration, task coordination, and result processing, while the chain component handles the blockchain-related operations including synchronization and reward distribution.

4.4 AIP Agent Interaction

Figure 7 presents the architectural framework for a blockchain-synchronized multi-agent communication protocol. The proposed system comprises four core components: distributed agents (agent0, agent1), a centralized agent hub, and a blockchain-based synchronization layer that maintains distributed state consistency. The protocol operates through two distinct phases: registration and interaction. In the registration phase, each agent initiates connection establishment by registering its identity with the agent hub and subsequently establishing a persistent communication channel. Concurrently, the agent hub synchronizes its local state with the blockchain to ensure global consistency across the distributed network. The interaction phase implements a request-response communication pattern with integrated authentication and authorization mechanisms. When agent0 requires communication with agent1, it first acquires authentication credentials specific to the target agent through the hub’s credential management system. Upon credential acquisition, agent0 transmits its request to the hub, which performs authentication validation by querying the blockchain-synchronized state. Following successful authentication, the hub redirects the request to agent1. The response flow follows a symmetric pattern: agent1’s response undergoes authentication verification before being redirected back to the requesting agent.

4.5 ERC-8004-Compatible Trust Layer

While AIP provides transport-level interoperability for agent-to-agent and agent-to-tool communication, it deliberately leaves the problem of cross-organizational trust to higher layers. To address this gap and align with the broader Ethereum agent ecosystem, Unibase extends AIP with native support for ERC-8004, the recently proposed “Trustless Agents” standard.¹ ERC-8004 defines three lightweight on-chain registries that allow agents to be discovered, evaluated, and validated across organizational boundaries without pre-existing trust, complementing transport protocols such as Google’s A2A and Anthropic’s MCP.²

Identity Registry Integration. Membase already assigns each agent a globally unique `MEMBASE_ID` and maintains its configuration and history.³ To remain compatible with ERC-8004, Unibase maps `MEMBASE_ID` entries to ERC-721-based agent identifiers in the ERC-8004 Identity Registry.⁴ For each registered agent, Unibase can optionally mint an ERC-8004-compliant Identity token on Ethereum or other EVM-compatible rollups, with the token URI referring to an agent registration file stored on UnibaseDA or compatible storage backends. This file includes AIP endpoints, supported protocols (e.g., gRPC, MCP, A2A), and additional metadata, allowing ERC-8004 indexers and wallets to treat Unibase agents as first-class, browsable on-chain entities.

Reputation Registry Integration. AIP’s task coordination layer (Section 4.2) already records task creation, assignment, and results on-chain.⁵ Unibase extends this flow by optionally emitting ERC-8004 Reputation Registry updates each time a task completes successfully. After a task finishes and payment is settled, the client or calling agent may submit standardized feedback entries (e.g., scores in $[0, 100]$, tags, and URIs pointing to detailed off-chain feedback) to the ERC-8004 Reputation Registry.⁶ This design allows external reputation services, dashboards, or meta-agents to consume a unified feed of agent feedback across ecosystems, while Unibase continues to store rich interaction traces and context in Membase and UnibaseDA.

Validation Registry Integration. UnibaseDA and the verification layer (Section 5) already support fraud proofs, zero-knowledge verification, and high-throughput data availability for AI workloads.⁷ These primitives align naturally with ERC-8004’s Validation Registry, which defines a generic interface for independent validation services to attest to an agent’s behavior using crypto-economic staking, zkML proofs, or TEE-based attestations.⁸ Unibase exposes contract hooks that allow external validators or on-chain coordinators to (1) request validation of specific tasks or agent behaviors, and (2) record validation outcomes in the ERC-8004 Validation Registry while storing supporting evidence—proofs, logs, or model traces—in UnibaseDA. This yields a spectrum of trust models, from lightweight reputation-only agents to high-assurance agents backed by cryptographic verification.

Benefits for the Open Agent Internet. By aligning AIP and Membase with ERC-8004, Unibase agents become natively interoperable with the broader Ethereum agent economy. Agents deployed on Unibase can be discovered through ERC-8004 Identity Registries, evaluated through shared Reputation and Validation registries, and still leverage Unibase’s decentralized memory and high-throughput data availability layer for their internal state. Conversely, third-party ERC-8004-compliant agents can integrate with Unibase as external services, using AIP for communication and UnibaseDA as a verifiable storage backend. This tight coupling between Unibase’s agent infrastructure and ERC-8004’s trust layer advances the vision of the Open Agent Internet by combining transport interoperability, persistent memory, and portable on-chain trust primitives in a single cohesive stack.

¹See ERC-8004: Trustless Agents [27], which introduces on-chain Identity, Reputation, and Validation registries for autonomous agents.

²ERC-8004 specification and ecosystem overviews are provided in [14, 30].

³See Section 3 for the dual-layer identity design of Membase.

⁴ERC-8004 builds its Identity Registry on ERC-721 with URISStorage, where each token’s URI points to an off-chain agent registration file [14, 27].

⁵See Figure 6 for the agent task execution workflow.

⁶The ERC-8004 Reputation Registry standardizes a shared interface for feedback entries while leaving aggregation and scoring logic to off-chain services [14, 30].

⁷See Section 5 for UnibaseDA’s on-chain verification, honest-one security model, and ZK integration.

⁸ERC-8004’s Validation Registry is designed to accept pluggable validation mechanisms, including TEE and zk-based approaches [30, 31].

4.6 x402–Native Payment Support

While AIP provides the communication and coordination fabric for autonomous agents, Unibase also requires a native mechanism for agents to pay for API calls, data, and services in a fully programmatic manner. To address this, Unibase integrates first-class support for the x402 payment standard, an open, chain-agnostic protocol that operationalizes the HTTP 402 “Payment Required” status code to enable native, crypto-based payments directly over HTTP.

In the x402 model, when an agent requests a paid resource, the server may respond with an HTTP 402 status containing machine-readable payment terms (price, asset, receiver, expiry). The agent then constructs and authorizes the required blockchain transaction and resubmits the request with an attached payment proof. Upon verification, the server returns the requested resource. This simple interaction pattern allows any HTTP endpoint to function as a pay-per-use “vending machine” for APIs, inference, data feeds, or tools, without accounts, sessions, or traditional billing infrastructure.

AIP–x402 Integration. x402-enabled endpoints allow any AIP-compatible service, including Membase and external tools, to define price schedules for specific routes. When an unpaid request is received, the service responds with an HTTP 402 message containing the full x402 payment metadata. Agent-side payment orchestration is handled by AIP client libraries, enabling agents to parse 402 responses, construct the required payment transactions (such as EVM transfers or permit-style authorizations), and resubmit the original request with a valid payment proof attached. Memory integration with Membase ensures that payment receipts, pricing data, and settlement outcomes can be stored as part of each agent’s long-term memory. This allows agents to make cost-aware decisions, optimize spending, and build verifiable reputations based on successful payment execution.

BNB Chain and Facilitators. Since Unibase initially deploys on BNB Chain, AIP integrates with x402 facilitators that support settlement on BNB-compatible networks. These facilitators generate payment requests, verify signatures, and finalize settlements on-chain. Agents can use x402 to pay for external APIs, data feeds, compute tasks, or other services using stablecoins or supported assets. AIP endpoints may also expose themselves as x402-protected services, allowing autonomous agents to operate as revenue-generating microservices without additional infrastructure.

All payment behavior remains anchored on BNB Chain, while interaction logs, records, and trust signals are stored in Membase and validated via the ERC-8004 trust layer.

AI-Native Commerce. x402 support tightly aligns Unibase with the broader vision of the Open Agent Internet. Agents can autonomously purchase context, tools, or services at fine-grained, per-request intervals. Agents may pay small amounts of stablecoins to retrieve research documents, run specialized inference queries, or access premium trading signals, all through standard HTTP interactions mediated by AIP. Payment receipts can be stored in Membase and, when necessary, committed to UnibaseDA for enhanced auditability. This design creates a unified substrate where communication, memory, trust, and payments work seamlessly together.

5 UnibaseDA: AI-Native Decentralized Storage & Data Availability

The exponential growth of AI applications in decentralized environments has exposed critical limitations in existing data availability (DA) solutions. Traditional Web3 storage systems, while adequate for conventional blockchain applications, fall short when confronted with the intensive storage demands of AI agents and models that require real-time access to massive datasets, context windows, and model parameters. Existing Data Availability Sampling (DAS) based architectures, including Celestia, EigenDA, and Avail, suffer from weakened security guarantees that rely on consensus protocols often less robust than Layer 1 chains, coupled with high transmission overhead that significantly impacts performance under large-scale data workloads.

To address these fundamental limitations, this chapter presents UnibaseDA, an AI-native decentralized data availability and storage layer that provides on-chain verifiability, ultra-high throughput, and massive scalability. UnibaseDA introduces a novel architecture that combines zero-knowledge proof verification, fraud-proof mechanisms, and optimized encoding strategies to achieve up to 100GB/s data throughput while maintaining security guarantees that require only

one honest validator, substantially improving upon the honest majority assumptions of existing solutions. UnibaseDA layer provides scalable, tamper-proof storage for operational data while maintaining cryptographic links to on-chain verification data. The DA layer implements advanced data availability proofs that ensure data integrity without requiring full replication across all network nodes. Regular synchronization between the Chain and DA layers ensures consistency while enabling the system to operate efficiently under varying network conditions. The DA layer supports multiple storage backends and implements intelligent data placement strategies to optimize access patterns for different types of agent data.

Features	Filecoin	Arweave	Storj	Greenfield	Unibase Storage
Capacity	EB+	~100TB	PB+	-	EB+
Data Availability	Low	Middle	Middle	Middle	HIGH
Access Control	N	N	N	Y	Y
Programmable	N	N	N	N	Y
Price (TB/Month)	LOW (0.19\$)	HIGH (30000\$)	LOW (5\$)	Middle (24\$)	LOW (5\$)
Use Cases	Cold data, high read latency	Small static file	Big Files	Decentralized terminal + centralized sp storage	High concurrency low-latency

Table 2: Comparison of Decentralized Storage Solutions

5.1 Decentralized Storage Infrastructure

UnibaseDA implements a hierarchical storage architecture where users submit blob commitments and Reed-Solomon (RS) parameters directly to the blockchain. The system partitions data into encoded fragments distributed across a network of storage nodes, supporting variable blob sizes and configurable lifespans to accommodate diverse storage requirements. This design enables exabyte-scale storage capacity with horizontal scalability supporting millions of storage nodes, while also providing custom private storage pools for enterprise deployments requiring dedicated environments.

The storage layer integrates seamlessly with the broader Unibase ecosystem, connecting to Unibase’s AI Native Storage Layer for verifiable, on-chain-accessible storage operations. Unibase supports programmable ownership and access control through smart contracts, enabling customizable lifecycle rules, data governance policies, and automated storage logic, including auto-expiry mechanisms, dynamic pricing, and access triggers tailored specifically for AI workflows.

5.2 On-Chain Verification

Unlike traditional DAS solutions that rely on off-chain verification, UnibaseDA implements comprehensive on-chain verification through two primary proof mechanisms: Encode Proofs and Duality Proofs. Encode Proofs verifies Reed-Solomon encoding correctness directly on-chain, ensuring that storage nodes correctly maintain their assigned encoded shards. Duality Proofs demonstrate continuous data availability throughout the committed validity window by requiring periodic proof submissions from storage nodes.

The verification system employs an optimistic model where proofs are assumed valid unless challenged, significantly reducing on-chain gas consumption while maintaining security guarantees. When missing or incorrect proofs are detected, any participant can verify them off-chain and initiate on-chain challenges. This approach, combined with proof aggregation techniques that batch multiple verification operations per storage node, achieves cost-effective validation without compromising security.

5.3 Security Model and Verification

5.3.1 Honest-One Security Guarantee

UnibaseDA implements a security model requiring only one honest validator to ensure system integrity, representing a significant improvement over honest majority assumptions employed by competing solutions. This security model is achieved through the combination of on-chain verification mechanisms and fraud-proof systems that enable any participant to challenge invalid operations.

The system eliminates the need for Data Availability Committees (DACs) by enabling any storage node to submit on-chain proofs. This decentralized approach to verification reduces single points of failure while maintaining strong security guarantees comparable to Layer 1 blockchain security models.

5.3.2 Zero-Knowledge Proof Integration

All data access and write operations are verified using zero-knowledge proofs, ensuring authenticity and auditability without compromising privacy. The ZK-based verification system enables tamper-proof storage operations while minimizing on-chain computational overhead. This integration supports compatibility with smart contract hooks and programmable storage lifecycles, enabling sophisticated AI applications to implement custom verification logic.

5.4 Storage Specification

Table 2 compares several decentralized storage solutions across various key metrics, including capacity expansion, data availability, access control, programmability, pricing, and use cases. Unibase Storage distinguishes itself in the decentralized storage market through four key competitive advantages that address critical gaps left by existing solutions. Unlike Filecoin’s [3] low availability issues or the moderate availability offered by Arweave [1], Storj [4], and Greenfield [2], Unibase Storage ensures high availability that makes it ideal for applications requiring real-time data access. Most significantly, Unibase is the only platform that supports both access control and programmability, making it substantially more versatile than competitors and enabling secure, customized, and automated storage solutions that cater to modern AI, Web3, and enterprise needs. From a cost perspective, Unibase offers exceptional value at 5/TB/month, matching Storj’s competitive pricing while delivering superior availability and performance—a stark contrast to Arweave’s prohibitive 30,000/TB price point that makes it impractical for large-scale deployments. Finally, while Filecoin and Storj are primarily optimized for cold storage and large file storage, Unibase excels in high-concurrency and low-latency environments, positioning it as the ideal choice for AI-driven applications, decentralized computing, and Web3 infrastructure that demand responsive, scalable storage solutions.

5.5 Comparison with Existing DA Solutions

Table 5.5 compares UnibaseDA, our proposed data availability solution, against five existing alternatives across key technical dimensions. UnibaseDA distinguishes itself through on-chain data verification on Ethereum [26], providing stronger security guarantees compared to the off-chain verification used by EigenDA [23, 25], Celestia [11, 17], Avail [13, 33], and 0G [9, 22]. Our solution employs an “Honest One” security model, requiring only a single honest participant to ensure data availability, which is more robust than the “Honest Majority” assumption (requiring more than 50% honest validators) used by all competing solutions. UnibaseDA achieves industry-leading throughput of 100 GBPS, significantly outperforming Ethereum’s 0.083 MBPS [16], EigenDA’s 10 MBPS [24], Celestia’s 1.5-10 MBPS [18], and doubling 0G’s 50 GBPS [10]. Like Celestia and 0G, UnibaseDA uses Fraud Proofs for encoding verification [21], which are more computationally efficient than the Validity Proofs used by EigenDA and Avail [32], while still maintaining security through challenge periods. UnibaseDA supports Data Availability Sampling (DAS) [15, 19], enabling light clients to verify data availability without downloading complete datasets, a critical feature absent in both Ethereum and EigenDA. Finally, UnibaseDA settles on Ethereum, leveraging its established security and decentralization [20] while dramatically improving throughput and trust assumptions for next-generation rollups and Layer 2 solutions.

Table 3: Comparison of Data Availability Solutions

Feature	Ethereum	EigenDA	Celestia	Avail	0G	Unibase DA
Data Verification	-	Off-chain	Off-chain	Off-chain	Off-chain	On-chain
Data Security	Honest Maj.	Honest Maj.	Honest Maj.	Honest Maj.	Honest Maj.	Honest One
Throughput	0.083 MBPS	10 MBPS	1.5-10 MBPS	-	50 GBPS	100 GBPS
Consensus	Ghost/Casper	None	Tendermint	BABE/GRANDPA	Tendermint	None
Support DAS	No	No	Yes	Yes	Yes	Yes
Encoding Proof	Validity	Validity	Fraud	Validity	Validity	Fraud
DA Settlement	Ethereum	Ethereum	Celestia	Avail	0G Chain	Ethereum

5.6 Conclusion

UnibaseDA represents a significant advancement in decentralized data availability infrastructure, specifically designed to meet the demanding requirements of AI-native applications. Through the combination of on-chain verification, honest-one security guarantees, and ultra-high throughput capabilities, the system addresses fundamental limitations of existing DAS-based architectures while providing the scalability and performance characteristics required for next-generation decentralized AI applications.

6 Roadmap




Figure 8: Unibase development roadmap: Four-phase deployment strategy from mainnet launch through full cross-platform interoperability.

Figure 8 shows Unibase Development Roadmap. Four-phase deployment strategy from mainnet launch through full cross-platform interoperability. The detailed deliverables are described below.

6.1 Phase 1: Foundation (September 2025)

Mainnet Launch on BNB Chain

The initial deployment of Unibase establishes the foundational infrastructure for the Open Agent Internet. This phase introduces the core capabilities that enable autonomous agents to operate with persistent identity and verifiable interactions on BNB Chain.

Immortal Agent Creation. Launch of the Membase identity registration system, enabling agents to establish permanent, cryptographically-secured identities on-chain. Each agent receives a globally unique MEMBASE_ID that persists across platforms and sessions, providing the foundation for long-term agent autonomy.

On-Chain Trading Infrastructure. Deployment of smart contract frameworks supporting agent-to-agent economic coordination and value exchange. This enables agents to participate in decentralized markets, execute transactions autonomously, and establish reputation through verifiable on-chain activity.

Basic Memory Hub Operations. Initial implementation of conversation memory storage and configuration management through the Agent Hub, supporting essential agent state persistence and cross-session continuity.

6.2 Interoperability

This phase focuses on expanding Unibase into a fully cross-chain, cross-platform coordination layer. Alongside AIP 2.0 and multi-chain memory synchronization, Unibase will deploy ERC-8004-compliant Identity, Reputation, and Validation registries across BNB Chain and selected Ethereum L2 networks. These registries will connect Unibase agents to the emerging ERC-8004 trust layer, enabling trustless interoperability with external agents and decentralized agent marketplaces.

6.3 Phase 2: Economic Coordination (October 2025)

Governance and Incentive Mechanisms

Building upon the mainnet foundation, this phase introduces sophisticated economic coordination mechanisms that align agent and stakeholder interests while enabling community-driven platform evolution.

veToken Staking System. Launch of vote-escrowed token mechanics that enable long-term stakeholders to participate in platform governance. Token holders can lock their holdings to receive voting power proportional to both stake amount and lock duration, aligning economic incentives with long-term network health.

Bribe Pool Mechanism. Implementation of incentive markets where participants can influence agent behavior and resource allocation through token-based coordination. This creates efficient mechanisms for directing agent activity toward high-value tasks and services.

Agent Governance Framework. Deployment of on-chain governance contracts enabling decentralized decision-making for protocol parameters, feature prioritization, and resource allocation. Agent owners and token holders participate in shaping platform evolution through transparent, verifiable voting mechanisms.

Interoperability. Deploy ERC-8004-compliant Identity, Reputation, and Validation registries across BNB Chain. These registries will connect Unibase agents to the emerging ERC-8004 trust layer, enabling trustless interoperability with external agents and decentralized agent marketplaces.

6.4 Phase 3: Verification Layer (Q4 2025)

On-Chain Memory Verification on BNB Chain

This critical phase activates the cryptographic verification infrastructure that ensures data integrity and enables trustless agent interactions at scale. The implementation of zero-knowledge proofs provides privacy-preserving verification without compromising performance.

ZK-Backed Membase Layer. Full activation of zero-knowledge proof systems for verifying agent memory operations on-chain. This enables agents to prove data authenticity and state transitions without revealing sensitive information, maintaining privacy while ensuring verifiability.

On-Chain Verification Contracts. Deployment of smart contracts that validate memory commitments, authorization proofs, and data availability guarantees. These contracts provide cryptographic assurance that agent memory remains tamper-proof and continuously accessible.

UnibaseDA Integration. Complete integration of the decentralized storage layer with on-chain verification mechanisms, enabling high-throughput data operations (up to 100 GB/s) while maintaining cryptographic links between off-chain storage and on-chain proofs. The fraud proof system allows any participant to dispute invalid memory operations or authorization claims, ensuring system integrity through economic incentives rather than trust assumptions.

6.5 Phase 4: Interoperability (Q1 2026)

AIP 2.0: Cross-Platform Memory Synchronization

The final phase of this roadmap realizes the vision of the Open Agent Internet by enabling seamless agent interoperability across diverse platforms and ecosystems. AIP 2.0 establishes universal standards for agent communication and memory sharing.

Cross-Platform Memory Sync. Launch of advanced synchronization protocols enabling agents to maintain consistent memory state across multiple platforms, blockchains, and execution environments. Agents can seamlessly migrate between ecosystems while preserving complete conversation history, knowledge bases, and operational context.

Enhanced Agent Discovery and Authentication. Implementation of decentralized registry systems enabling agents to discover and establish connections with complementary agents across platforms. This supports dynamic multi-agent workflows where capabilities can be automatically matched to task requirements. Deployment of cross-chain identity verification mechanisms allows agents to authenticate and authorize interactions regardless of the underlying blockchain or platform, eliminating integration friction while maintaining strong security guarantees.

Protocol Bridge Enhancements. Expanded compatibility layers ensuring seamless interoperability with existing agent frameworks including ElizaOS, Virtuals Protocol, and Swarms Framework. These MCP and gRPC bridges enable legacy agents to leverage Unibase’s memory and identity infrastructure without requiring complete rewrites.

Multi-Chain Task Coordination. Extension of the blockchain task coordination system to support cross-chain workflows, enabling agents to participate in tasks that span multiple blockchain ecosystems while maintaining unified state management and reward distribution.

6.6 Beyond Q1 2026: Future Directions

Following the completion of these foundational phases, the Unibase ecosystem will continue evolving through community-driven development and integration with emerging AI technologies. Future directions include scaling to support millions of concurrent agents, implementing advanced reputation systems for agent trust management, and expanding the UnibaseDA infrastructure to achieve exabyte-scale storage capacity with sub-100ms latency guarantees. The platform will also explore integration with emerging AI model architectures, enhanced privacy-preserving computation capabilities through advanced zero-knowledge techniques, and sophisticated multi-agent coordination protocols that enable increasingly complex autonomous workflows across the Open Agent Internet.

7 Ecosystem

7.1 Ecosystem Applications and Use Cases

BitAgent [6] Platform: BitAgent represents a comprehensive multi-agent orchestration platform that showcases Unibase’s advanced capabilities for managing complex agent coordination and autonomous interaction. The platform streamlines agent deployment through simplified creation processes and robust configuration management systems, while supporting agent staking mechanisms that enable economic participation in the network. It enables collaborative workflows where multiple agents can execute tasks together, automatically matching capabilities to optimize performance. The system incorporates economic coordination mechanisms using token-based incentive structures that align agent motivations and ensure effective collaboration across the network, facilitating seamless launching and operation of autonomous agents at scale.

TwinX [8] Social Agent Builder: TwinX offers rapid automated agent creation, transforming users’ tweets into fully functional self-learning AI agents with persistent memory capabilities in just five minutes. The platform performs intelligent content analysis, parsing and modeling user communication patterns to understand individual expression styles and preferences. Through advanced personality synthesis capabilities, it generates AI-driven personality models that accurately reflect users’ social interactions and behavioral tendencies. The system maintains continuous learning functionality with built-in memory systems, adapting and refining agent personalities based on

ongoing user feedback and real-world interactions to ensure authentic representation and evolving intelligence over time.

TradingFlow [7] Autonomous Strategy Engine: TradingFlow operates as an AI-native autonomous trading agent that revolutionizes strategy development through natural language interface capabilities. The platform excels at strategy interpretation, converting complex natural language descriptions into executable trading logic without requiring technical programming knowledge, enabling traders to articulate strategies conversationally and have them automatically implemented. It incorporates sophisticated risk management features that automatically assess portfolio risk levels and determine optimal position sizing for each trade. The system provides comprehensive performance attribution analysis, delivering detailed insights into strategy effectiveness while offering data-driven recommendations for continuous improvement and optimization of autonomous trading operations.

Beeper Intent Agent: Beeper serves as an intelligent intent-driven agent specifically designed for social crypto interactions on Twitter. The platform enables seamless crypto tipping functionality, allowing users to reward content creators and engage in social transactions directly through natural language commands. It facilitates the distribution and claiming of red packets, bringing traditional gifting mechanisms to the crypto ecosystem with automated coordination. The agent also provides streamlined access to DeFi actions, interpreting user intent to execute complex decentralized finance operations through simple social media interactions, thereby bridging the gap between social engagement and blockchain-based financial activities.

7.2 Ecosystem Patners

Unibase maintains active integrations with leading agent frameworks to enable seamless deployment and coordination of AI agents across diverse environments. Through its integration with ElizaOS [34], Unibase provides a robust memory backend for the modular AI runtime environment, ensuring persistent state management and knowledge retention across agent interactions. The Virtuals Protocol [5] integration enables composable agent deployment with persistent identity management, allowing agents to maintain consistent personas and capabilities across different contexts and platforms. Additionally, the Swarms Framework [28] integration facilitates multi-agent task orchestration with decentralized coordination, enabling complex collaborative workflows where multiple agents can work together autonomously while leveraging Unibase’s unified data layer for shared knowledge and coordination signals.

8 Economics and Tokenomics

8.1 Overview

The Unibase ecosystem is powered by the UB token, which serves as the fundamental economic primitive enabling decentralized coordination, incentive alignment, and sustainable network growth. Unlike traditional utility tokens that serve singular purposes, UB is designed as a multi-dimensional coordination mechanism that addresses the unique economic challenges of decentralized AI infrastructure. The token simultaneously functions as a medium of exchange for protocol services, a governance instrument for collective decision-making, a staking asset for infrastructure provision, and an incentive mechanism for knowledge contribution. This integrated economic design ensures that all stakeholders—from agent developers and infrastructure operators to knowledge contributors and governance participants—are aligned toward the shared objective of building a robust, scalable, and open AI memory layer.

The UB token economy is architected around four primary utility pillars that collectively address the economic requirements of a decentralized AI infrastructure: Protocol Fees, Governance Rights, Infrastructure Staking, and Knowledge Mining Rewards. Each pillar serves distinct but complementary roles in sustaining network operations, ensuring quality standards, and incentivizing long-term participation. This multifaceted approach creates self-reinforcing network effects where increased agent deployment drives demand for memory services, which in turn incentivizes infrastructure expansion and knowledge contribution, ultimately strengthening governance participation and protocol resilience.

8.2 Token Distribution and Vesting Schedule

The UB token has a total supply of **10 billion tokens** distributed across seven categories designed to balance community ownership, ecosystem growth, team incentives, and market liquidity. The

distribution incorporates carefully structured vesting schedules to ensure long-term alignment and prevent excessive market pressure.

Category	Allocation	Token Amount	TGE Unlock	Vesting Schedule
Community	35%	3,500,000,000	5%	6-month cliff, then 24-month linear release
Ecosystem	10%	1,000,000,000	0%	6-month cliff, then 24-month linear release
Treasury	20%	2,000,000,000	0%	6-month cliff, then 24-month linear release
Team & Advisors	18%	1,800,000,000	0%	6-month cliff, then 24-month linear release
Marketing	10%	1,000,000,000	4.25%	Linear vesting over 4 months
Liquidity	5%	500,000,000	5%	Fully unlocked at TGE
Binance Alpha	2%	200,000,000	2%	Fully unlocked at TGE

Table 4: UB Token Distribution and Vesting Schedule (Total Supply: 10,000,000,000 UB)

8.3 Protocol Fees

The Unibase protocol requires UB tokens for fundamental operations across its three core infrastructure layers: agent deployment, memory storage, and interoperability services. This fee structure ensures sustainable protocol operations while creating organic demand proportional to network utilization.

Agent Deployment Fees: Deploying autonomous agents on the Unibase network requires UB token payment to register unique agent identities through Membase. These fees cover the computational costs of on-chain identity verification, cryptographic key generation, and initial configuration synchronization. The deployment fee structure is calibrated to balance accessibility for developers with sufficient economic weight to prevent identity spam and Sybil attacks. Fees scale based on agent complexity, with basic agents requiring minimal deployment costs while sophisticated multi-agent systems with enhanced coordination capabilities incur proportionally higher fees.

Memory Storage Fees: Persistent memory storage through Membase constitutes a core value proposition of Unibase, enabling agents to maintain context across sessions and platforms. Storage fees are denominated in UB and calculated based on data volume, retention duration, and access frequency requirements. The fee model distinguishes between different memory modalities—with conversation memory, knowledge base storage, and high-frequency operational data each priced according to their distinct performance characteristics. Storage fees incorporate both initial upload costs and ongoing maintenance fees that ensure data availability and integrity across the decentralized storage network. This dual-component pricing structure incentivizes efficient data management while supporting sustainable compensation for storage node operators.

AIP Protocol Access Fees: Agents utilizing the Agent Interoperability Protocol (AIP) for cross-platform communication, tool integration, and collaborative task execution incur protocol access fees denominated in UB. These fees support the infrastructure required for secure agent-to-agent messaging, authentication verification, and permission management. The AIP fee structure implements usage-based pricing that scales with interaction frequency and complexity, ensuring that high-throughput agent applications contribute proportionally to network maintenance while keeping costs minimal for lightweight interactions. Protocol fees also fund ongoing development of interoperability standards and integration with emerging agent frameworks.

x402-Based Fee Collection In addition to direct on-chain transactions, Unibase anticipates that a significant share of protocol fees will flow through x402-compatible payment rails. For AIP-protected APIs and Membase-backed services, protocol fees MAY be collected via x402 micro-payments denominated in UB or other supported assets, with settlement executed on BNB Chain by compatible facilitators. In this model, protocol endpoints expose HTTP 402 paywalls that quote prices and accepted assets, while Membase records payment metadata and UnibaseDA stores associated receipts or proofs when stronger auditability is required. This design aligns Unibase’s fee model with emerging, agent-native payment infrastructure while preserving the same security, governance, and incentive properties described in Sections 5 and 8.

8.4 Governance and veUB

Unibase implements a vote-escrowed governance model where UB token holders can lock their tokens to obtain veUB (vote-escrowed UB), granting enhanced governance rights and protocol influence. This time-locked voting mechanism ensures that governance participants have long-term alignment with protocol success, as voting power correlates with both token quantity and lock duration.

Lock Duration and Voting Power: Token holders can lock UB for periods ranging from one week to four years, with voting power scaling linearly with lock duration. Maximum-length locks (four years) grant 1:1 veUB conversion, while shorter locks receive proportionally reduced voting weights. This mechanism ensures that stakeholders making longer-term commitments exert greater influence over protocol evolution, aligning decision-making authority with sustained network participation. veUB balances decay linearly over time, requiring participants to extend their locks periodically to maintain voting power, creating natural governance turnover while rewarding sustained commitment.

Governance Scope: veUB holders participate in critical protocol decisions including parameter adjustments for fee structures, storage pricing models, and reward distribution mechanisms. Governance authority extends to treasury management, determining allocation of protocol revenues toward ecosystem development, research funding, and strategic initiatives. The governance framework also encompasses protocol upgrades, new feature proposals, and integration decisions with external agent frameworks and infrastructure providers. Emergency response protocols grant veUB holders authority to intervene in critical security situations, enabling rapid community-driven responses to potential vulnerabilities.

Reward Allocation: Beyond voting rights, veUB holders receive preferential allocation of protocol revenues generated from fees collected across agent deployment, memory storage, and AIP access. Revenue sharing creates direct economic alignment between governance participants and protocol growth, as increased network utilization translates to enhanced rewards for long-term token lockers. This revenue distribution mechanism functions as a decentralized dividend system where veUB holders effectively serve as protocol stakeholders, benefiting financially from prudent governance decisions that drive adoption and usage.

8.5 Memory Node Staking

The decentralized memory and data availability infrastructure underlying Membase and UnibaseDA requires a distributed network of high-performance storage nodes operated by independent participants. UB token staking serves as both an economic commitment mechanism and a quality assurance system, ensuring that infrastructure operators maintain reliable service standards.

Node Operator Requirements: Operating a memory node or DA node requires staking a minimum threshold of UB tokens as collateral, with stake amounts scaling based on storage capacity commitments and performance tier selections. Node operators must demonstrate technical capabilities including minimum bandwidth specifications, storage reliability guarantees, and uptime commitments. The staking requirement creates economic accountability, as operators risk stake slashing for service failures, extended downtime, or failure to maintain data availability guarantees. This economic security model ensures that infrastructure participants have strong incentives to maintain high-quality service standards.

Horizontal Scalability: The staking model is explicitly designed to support massive horizontal scaling, accommodating growth from hundreds to potentially millions of storage nodes as network demands increase. Unlike traditional blockchain validators that face inherent scaling limitations due to consensus requirements, Unibase memory nodes operate with greater independence, performing specialized storage and retrieval functions without requiring global state synchronization. This architecture enables linear scalability where network capacity grows proportionally with the number of active nodes, allowing Unibase to accommodate exponential increases in agent deployment and memory storage demands.

Reward Distribution: Staked node operators earn UB token rewards proportional to their service contributions, calculated based on storage capacity provided, data served, and performance metrics including retrieval latency and uptime reliability. The reward algorithm incorporates quality-of-service parameters, providing enhanced compensation for nodes that consistently exceed minimum requirements. This performance-based reward structure creates competitive dynamics that drive continuous infrastructure improvement, as operators optimize their systems to capture greater reward shares. Additionally, rewards scale with network utilization, ensuring that operators benefit directly from ecosystem growth and increased demand for memory services.

Slashing Mechanisms: To maintain service quality and data integrity, the staking system implements automated slashing penalties for node operators who fail to meet commitments. Slashing conditions include extended downtime exceeding specified thresholds, data loss or corruption, failed availability proofs, and fraudulent behavior in proof submissions. Slash amounts scale with violation severity, ranging from minor penalties for transient failures to complete stake forfeiture for critical security breaches or malicious behavior. These economic penalties, combined with reputation scoring systems, create strong deterrents against poor service quality while maintaining fair treatment for operators experiencing legitimate technical difficulties.

8.6 Knowledge Mining

Beyond infrastructure provision, Unibase introduces a novel economic mechanism called Knowledge Mining that rewards contributions to the open memory layer. This system recognizes that high-quality prompts, curated memories, and reusable knowledge bases constitute valuable public goods that benefit the entire agent ecosystem.

Contribution Types: Knowledge mining rewards encompass multiple contribution categories, each valued according to their utility and impact. Prompt contributions include well-structured agent instructions, task templates, and workflow definitions that can be reused across multiple agent deployments. Memory contributions consist of curated conversation histories, interaction patterns, and domain-specific knowledge that can accelerate new agent training and improve reasoning capabilities. Knowledge base contributions include structured datasets, domain expertise documentation, and semantic relationship mappings that enhance agent understanding and decision-making across specialized domains.

Quality Assessment: To prevent low-quality submissions and ensure knowledge contributions maintain high standards, the system implements multi-layered quality assessment mechanisms. Initial filtering applies automated metrics evaluating contribution completeness, structural validity, and metadata quality. Secondary assessment involves peer review by experienced community members who evaluate practical utility, accuracy, and reusability potential. Top-tier contributions undergo expert validation for specialized domains requiring technical verification. This tiered quality system ensures that rewards correlate with genuine value creation rather than volume-based gaming.

Reward Mechanisms: Knowledge mining rewards are distributed from a dedicated allocation pool, with payout amounts determined by contribution quality scores, usage frequency, and long-term impact metrics. High-quality contributions that see sustained utilization across multiple agents and applications receive enhanced rewards over time, creating incentives for lasting value creation rather than short-term gaming. The reward structure also accounts for contribution novelty, providing premium compensation for pioneering work in underserved domains or innovative prompt engineering techniques that advance the state of agent capabilities.

Ecosystem Effects: By monetizing knowledge contribution, Unibase creates economic incentives that accelerate the development of public agent resources, reducing barriers to entry for new developers and improving overall agent quality across the network. This economic model transforms knowledge sharing from an altruistic activity into a sustainable profession, potentially attracting specialized domain experts, prompt engineers, and AI researchers to contribute their expertise to the open ecosystem. Over time, knowledge mining is expected to generate comprehensive libraries of reusable agent components, accelerating innovation cycles and enabling more sophisticated agent applications through compositional knowledge assembly.

8.7 Economic Sustainability and Network Effects

The integrated UB token utility model creates self-reinforcing economic dynamics that support long-term protocol sustainability and growth. Protocol fee revenues fund ongoing development, security audits, and ecosystem grants while providing rewards to governance participants and infrastructure operators. As agent deployment increases, demand for memory storage grows, driving up node operator rewards and attracting additional infrastructure capacity. Enhanced infrastructure performance supports more sophisticated agent applications, expanding the addressable market and driving further protocol adoption.

Knowledge mining contributions improve agent capabilities and reduce development costs, accelerating new project launches and increasing protocol utilization. This positive feedback loop between infrastructure provision, knowledge contribution, and agent deployment creates robust network effects where ecosystem value compounds over time. The governance mechanism ensures that economic parameters remain calibrated to market conditions, with veUB holders adjusting

fee structures, reward distributions, and staking requirements to maintain competitive positioning while ensuring protocol sustainability.

Unlike purely speculative token models, UB’s multi-dimensional utility ensures that token demand correlates directly with real network usage, creating organic price support from fundamental value creation. As the agent ecosystem matures and Unibase establishes itself as critical infrastructure for autonomous AI systems, the combination of protocol fees, governance rights, infrastructure rewards, and knowledge mining incentives positions UB as an essential coordination mechanism for decentralized AI development, with economic value accruing to long-term participants who contribute to protocol growth and ecosystem expansion.

9 Conclusion

Unibase represents a foundational breakthrough in decentralized AI infrastructure, addressing the critical limitations that have constrained autonomous agent development and deployment at scale. Through the seamless integration of Membase’s dual-layer identity system and hierarchical memory management, AIP’s Web3-native communication protocols, and UnibaseDA’s high-throughput verifiable storage layer, the platform establishes the technical foundation necessary for the Open Agent Internet vision. The system’s ability to maintain persistent agent memory across platforms, enable trustless inter-agent collaboration through cryptographic verification, and provide sub-100ms data access at exabyte scale positions it as essential infrastructure for the next generation of decentralized AI applications. As demonstrated through early ecosystem partners like BitAgent, TwinX, and TradingFlow, Unibase’s modular architecture not only supports diverse AI use cases but also enables new forms of agent coordination and economic interaction that were previously impossible in centralized systems. By solving the fundamental challenges of memory persistence, platform interoperability, and data sovereignty while maintaining the performance characteristics required for real-time AI workloads, Unibase paves the way for a truly autonomous, decentralized, and composable agent ecosystem that preserves user control and promotes innovation at global scale.

References

- [1] Arweave 2024. *Arweave: A Community-Driven Ecosystem*. Arweave. <https://arweave.org/> Official website of the Arweave permanent storage network and permaweb ecosystem.
- [2] BNB Chain 2024. *BNB Greenfield: A Decentralized Data Storage System and Economy*. BNB Chain. <https://greenfield.bnbchain.org/en> Official website of BNB Greenfield decentralized storage network.
- [3] Filecoin Foundation 2024. *Filecoin: A Decentralized Storage Network for the World’s Information*. Filecoin Foundation. <https://filecoin.io/> Official website of the Filecoin decentralized storage network.
- [4] Storj Labs 2024. *Storj: Smarter Cloud Storage for Your Business*. Storj Labs. <https://www.storj.io/> Official website of the Storj decentralized cloud storage platform.
- [5] Virtuals Protocol 2024. *Virtuals Protocol: Society of AI Agents*. Virtuals Protocol. <https://app.virtuals.io/> Composable agent deployment platform with persistent identity and tokenized AI agents.
- [6] BitAgent 2025. *Bitagent - Kickstart the Era of AI Knowledge Mining*. BitAgent. <https://www.bitagent.io/> First inclusive AI Agent platform on the entire chain for AI knowledge mining.
- [7] TradingFlow 2025. *TradingFlow - AI Powered Financial Market Research*. TradingFlow. <https://tradingflow.ai/> AI-powered platform for financial market research and analysis.
- [8] TwinX 2025. *TwinX App - Build Your AI Twin in 5 Minutes from Your Posts*. TwinX. <https://www.twinx.fun/> Self-evolving Agent platform for building personalized AI agents from social content.
- [9] 0G Labs. 2025. *0G’s Data Availability Layer*. 0G Blog. <https://0g.ai/blog/0g-s-data-availability-layer>
- [10] 0G Labs. 2025. *The Utility of 0G*. 0G Blog. <https://0g.ai/blog/the-utility-of-0g>

- [11] Mustafa Al-Bassam, John Adler, and Ismail Khoffi. 2022. *Celestia: A Scalable General-Purpose Data Availability Layer for Decentralized Apps and Trust-minimized Sidechains*. Celestia Blog. <https://blog.celestia.org/celestia-a-scalable-general-purpose-data-availability-layer-for-decentralized-apps-and-trust-minimized-sidechains>.
- [12] Anthropic. 2024. Introducing the Model Context Protocol. <https://www.anthropic.com/news/model-context-protocol> Open standard for connecting AI assistants to data sources and external tools.
- [13] Avail Project. 2025. *A Guide to Selecting the Right Data Availability Layer*. Avail Blog. <https://blog.availproject.org/a-guide-to-selecting-the-right-data-availability-layer/>
- [14] Backpack Exchange. 2025. *ERC-8004 Explained: Ethereum’s AI Agent Standard Guide*. No public URL available.
- [15] Blockworks Research. 2025. *Why data availability sampling matters for blockchain scaling*. Blockworks. <https://blockworks.co/news/data-availability-ethereum>
- [16] Vitalik Buterin. 2025. *Ethereum’s Fusaka Upgrade’s PeerDAS Feature to Enable Data Availability Scaling*. BanklessTimes. <https://www.banklesstimes.com/articles/2025/09/25/ethereums-fusaka-upgrades-peerdas-feature-to-enable-data-availability-scaling/>
- [17] Celestia Labs. 2024. *Celestia - The Modular Blockchain Network*. Celestia.org. <https://celestia.org/>
- [18] Celestia Labs. 2024. *Data Availability FAQ*. Celestia Documentation. <https://docs.celestia.org/learn/how-celestia-works/data-availability-faq>
- [19] Arunima Chaudhuri et al. 2024. On the Design of Ethereum Data Availability Sampling: A Comprehensive Simulation Study. *arXiv preprint arXiv:2407.18085* (7 2024). <https://arxiv.org/abs/2407.18085>
- [20] Coinbase. 2024. *Guide to EigenDA*. Coinbase Developer Platform. <https://www.coinbase.com/developer-platform/discover/protocol-guides/guide-to-eigenda>
- [21] CoinMarketCap. 2021. *What Is Data Availability?* CoinMarketCap Academy. <https://coinmarketcap.com/academy/article/what-is-data-availability>
- [22] Sami Dotan. 2025. *0G Labs raises \$35M for data availability, spurring competition with Celestia, EigenLayer, others*. Blockworks. <https://blockworks.co/news/0g-labs-fundraise-data-availability-layer>
- [23] Eigen Labs. 2024. *EigenDA Overview*. EigenDA Documentation. <https://docs.eigenda.xyz/overview>
- [24] Eigen Labs. 2024. *EigenDA: Secure, high-throughput, and decentralized Data Availability*. GitHub Repository. <https://github.com/Layr-Labs/eigenda>
- [25] Eigen Labs. 2024. *Intro to EigenDA: Hyperscale Data Availability for Rollups*. EigenCloud Blog. <https://blog.eigencloud.xyz/intro-to-eigenda-hyperscale-data-availability-for-rollups/>
- [26] Ethereum Foundation. 2024. *Data Availability*. ethereum.org. <https://ethereum.org/en/developers/docs/data-availability/>
- [27] Ethereum Foundation. 2025. *ERC-8004: Trustless Agents*. <https://eips.ethereum.org/EIPS/eip-8004>
- [28] Kye Gomez. 2024. *Swarms: The Enterprise-Grade Production-Ready Multi-Agent Orchestration Framework*. GitHub Repository. <https://github.com/kyegomez/swarms> Multi-agent orchestration framework with decentralized coordination and hierarchical agent swarms.
- [29] Google. 2025. *Announcing the Agent2Agent Protocol (A2A)*. <https://developers.googleblog.com/en/a2a-a-new-era-of-agent-interoperability/> Open protocol for secure agent-to-agent communication and collaboration.

- [30] Paweł Kuryłowicz. 2025. *ERC-8004: A Practical Explainer for Trustless Agents*. No public URL available.
- [31] Caer Lower. 2025. *ERC-8004: Building Trustless Autonomous Agents with TEEs*. No public URL available.
- [32] Gisele Schout. 2024. *Data Availability for Blockchains with Avail*. Medium - Stakin. <https://medium.com/stakin/data-availability-for-blockchains-with-polygon-avail-37b1c5f017ef>
- [33] Marc Schär. 2024. *Avail, Long-Awaited 'Data Availability' Blockchain Project, Launches Main Network*. CoinDesk. <https://www.coindesk.com/tech/2024/07/23/avail-long-awaited-blockchain-project-for-data-availability-launches-main-network>
- [34] ElizaOS Team. 2024. Eliza: Autonomous Agents for Everyone. GitHub Repository. <https://github.com/elizaOS/eliza> TypeScript framework for building AI agents with persistent personalities and cross-platform capabilities.